# An Introduction to the FAIR Controls Analytics Model (FAIR-CAM™)

# Table of Contents

# Introduction

*There's an old saying in marketing;*
*"Half of your marketing dollars are wasted. You just don't know which half."*

With that in mind, consider these questions:

1. Which cybersecurity control where you work provides the most value to the organization?
2. Which cybersecurity control provides the least value?
3. Would your answers be the same as someone else's from your organization?

The third question acts as a litmus test in the event that someone tries to answer either of the first two questions.

Let's imagine trying to answer the first two questions using any of the commonly applied control frameworks (e.g., NIST 800-53, PCI DSS, CIS Controls, COBIT, NIST CSF, etc.). An organization can evaluate its cybersecurity program against their framework of choice — determining whether each control in the framework is implemented and, if so, rating its condition in some fashion (e.g., using a 1-through-5 ordinal scale). At the end of the exercise, the organization would have a list of which controls were missing, and which scored highest, lowest, etc.[1]

There are at least three things to keep in mind though:

- *Since the value of any cybersecurity or risk management control boils down to how much it reduces risk,* we have to understand which loss event scenarios a control is relevant to, and how significantly the control affects the frequency or magnitude of those scenarios. This is not typically part of the evaluation process when cybersecurity programs are evaluated using common control or maturity frameworks, which means the value of each control isn't determined.
- Without knowing the risk-reduction value of its controls, an organization may inadvertently invest heavily in one or more controls that aren't particularly relevant to, or effective against, the risks it faces. When this is the case, the organization would have high scores for those less-relevant, less-valuable controls. For the same reason, the organization may under-invest in more important controls, which would result in lower scores for those controls. Organizations also sometimes invest more-or-less equally in as many controls as possible, which invariably results in under-investment in some controls and over-investment in others.
- All controls have relationships with, and dependencies upon, other controls, which is not accounted for in common control frameworks. As a result, weaknesses in some controls can diminish the efficacy of other controls. For example, the efficacy of an organization's patching process is highly dependent upon the efficacy of the organization's vulnerability identification capabilities, as well as its threat intelligence capabilities, and its risk analysis capabilities. If one or more of those capabilities is deficient, then the efficacy of patching will also be affected.

The bottom line is that simply scoring your organization's cybersecurity program based on common control or maturity framework doesn't provide meaningful insight into which controls are most or least valuable. And when organizations are unable to reliably understand the value they receive for their investments in risk management, then it's impossible to know whether they are overspending, underspending, or misallocating their resources.

Now, instead imagine that your organization could reliably answer the first two questions above, and that the answer to the third one was "Yes." Furthermore, imagine that the answers were the result of empirical

[1] Some frameworks focus on "control outcomes" rather than the controls themselves. However, with regard to answering the two value questions, nothing changes. For example, nothing in those frameworks helps us to know whether "Threats, both internal and external, are identified and documented" (NIST CSF ID.RA-3) is more valuable than "Physical access to assets is managed and protected" (NIST CSF PR.AC-2).

measurements of performance and efficacy that could be validated, instead of ordinal 1-through-5 scores. Your organization would have an entirely different level of clarity regarding when, where, and how it should apply its risk management resources, and could reliably justify those decisions to outside stakeholders.

The FAIR Controls Analytics Model (FAIR-CAM™) ontology described in this document is intended to enable this kind of focus and these capabilities.[2]

## How Is This Different?

*"In the 19th century we had a relatively advanced understanding of anatomy,*
*but we had a terrible understanding of physiology.*
*We knew what was happening, but we didn't know why it was happening."*

A RETIRED SURGEON

Human anatomy is essentially a list of the parts that make up a human body (e.g., the lungs, the heart, the nervous system, blood vessels, etc.). In contrast, human physiology describes how those parts perform their functions, both individually (e.g., the exchange of oxygen and carbon dioxide within the lungs) and as an overall system (e.g., how the nervous system senses levels of oxygen and carbon dioxide to stimulate respiration, and how other organs react to different levels of those gases).

Similar to a human body, the controls landscape functions as a set of complex, interdependent parts. As a result, we can only reliably and defensibly answer questions like those at the beginning of this paper if we understand "controls physiology." This focus on how the control landscape works is what the FAIR-CAM™ ontology provides. This will complement, rather than displace, existing risk management control frameworks (e.g., NIST CSF, HITRUST, COBIT, CIS, ISO2700x, etc.), which are descriptions of good practices and desirable control outcomes — i.e., controls anatomy. The FAIR-CAM™ ontology doesn't attempt to describe best practices, nor does it provide a list of cybersecurity outcomes like "Network integrity is protected" (NIST CSF PR.AC-5). Instead, it provides explicit descriptions of the control functions that affect risk, as well as the relationships and interdependencies between control functions.

For example, a typical control framework today might dictate that an education and awareness program is an important element within a cybersecurity program. And intuitively, we understand that this will help to reduce risk. But in order to understand how much risk it reduces within our organization (i.e., its value) so that we can reliably prioritize it, we have to understand several things:

- How education and awareness affects risk (it affects risk *indirectly, by improving the reliability of other controls*).
- Which controls are affected by an education and awareness curriculum (e.g., authentication password choices, personnel's ability to recognize phishing, etc.)
- Which loss event scenarios those controls are relevant to (e.g., compromise of user systems via phishing resulting in a ransomware outage, etc.)
- The condition of controls that education and awareness is dependent upon for its efficacy (e.g., risk appetite, policies, threat intelligence, etc.)
- As well as the condition of other controls that affect risk in those same scenarios (e.g., anti-malware technologies, etc., which also reduce risk associated with ransomware)

[2] Readers who are familiar with the the book Measuring and Managing Information Risk: A FAIR Approach will notice some differences between what is described in chapter 11 of that book versus FAIR-CAM. That chapter was written at a much earlier stage of my research into controls and, although the contents of that chapter were directionally correct, many details were missing and some were inaccurate, and it lacked any clear description of how to apply the controls model.

Highly experienced risk management professionals may have developed an "intuition" about controls that incorporate many of these considerations, however intuition can also be strongly affected by gaps in one's mental model, as well as by biases. Furthermore, accurate mental models tend to develop very slowly in complex problem spaces, and are difficult to explain, troubleshoot, or transfer.

The FAIR-CAM™ ontology documents how controls physiology works, which can refine and solidify the mental models of highly experienced professionals, and significantly accelerate an understanding of how the controls landscape works for newer professionals. It also provides a language for more clearly and concisely discussing controls. Lastly, it will enable accurate and reliable use of cybersecurity telemetry.

When combined with a scenario-based analysis model like FAIR, and well-defined "anatomy-like" control frameworks, the "control physiology" described by FAIR-CAM™ provides the means to reliably measure, analyze, forecast, and empirically validate control efficacy and value.

## Setting Expectations

*"Everything should be made as simple as possible, but not simpler."*

ALBERT EINSTEIN (OSTENSIBLY)

Modern medicine isn't complex because doctors and scientists made it unnecessarily complex. It's complex because human physiology and pathology are inherently complex. Fortunately, the principles and elements that comprise FAIR-CAM™ are less complex than those that underlie modern medicine. They are, however, more complex than how the cybersecurity profession has historically approached the controls landscape.

In order to embrace a more complex approach to controls, it's important to understand at least a few of the key reasons why the additional complexity is necessary and beneficial. With that in mind, here are a few of the contributing factors:

- **Units of measurement:** In order to ensure that control performance can be empirically measured and validated, each control function has to have an actual unit of measurement (e.g., frequency, probability, time, etc.). This differs from common control frameworks today that rely on abstract ordinal measurements (e.g., 1-through-5 scales, etc.). However, using explicit units of measurement also means that control functions need to be strictly differentiated — i.e., you can't munge related controls into a single high-level function as is often found in common frameworks today. For example, "Detection" has long been one of the commonly recognized control functions in risk management. And organizations typically might score their Detection capabilities as a "3" or a "2" based on the technologies and processes they have in place. But it's hard to know exactly what those scores mean, as they're simply ordinal labels, and it's impossible to reliably know how much additional risk reduction an organization gets from having a "3" level detection capability versus a "2". But detection of loss events actually is made up of three distinctly different and measurable components:

    1. The degree of visibility into activity that might be anomalous or malicious (e.g., what percentage of the network traffic is being captured)
    2. The elapsed time between reviews of data provided by visibility controls (e.g., how often does someone examine the captured network traffic), and
    3. The probability that anomalous or malicious activity will be recognized as a problem (e.g., being able to differentiate normal network traffic from abnormal activity)

Yes, evaluating these functions separately, measuring them empirically, and then combining those values to understand the efficacy and risk reduction value of an organization's detection capabilities is more involved than subjectively scoring Detection as a "3". But the increased decision-making utility and reliability of an empirically driven approach are significantly higher.

- **Dependencies and feedback loops:** A second contributing factor was that FAIR-CAM™ had to account for the dependencies and feedback loops that exist amongst controls within a risk management program. This is crucial for understanding how the control landscape actually works, and to enable reliable measurement of control performance and value.
- **Cybersecurity landscape complexity:** A third contributing factor is that loss events can play out over multiple layers of network, system, application, and user account architecture, which introduces myriad potential pathways for loss events. As a result, FAIR-CAM™ also needed to support the notion of control layers across this complex landscape.

Fortunately, as we gain insights from a deeper understanding of the controls landscape, we should be able to boil some aspects down to simpler, more operationally pragmatic heuristics without losing measurement validity. To use a medical analogy, this would be like creating more effective field medicine based on significant advancements in understanding human physiology.

We also should be able to automate a significant portion of controls analytics by ingesting cybersecurity telemetry and other available metrics to feed the model. In fact, this work has already begun to be developed.

Something else to keep in mind is that not everyone in the profession will need to know or work with FAIR-CAM™ at its deepest levels. Risk management will need its version of field medics, as well as physicians and physiologists — each with different depths of understanding and application.

As a final expectation-setting note, please recognize that this document only introduces the principles underlying the model, the model elements, and some distinct terminology. It does not include comprehensive examples of its application or a description of the mathematical formulas that enable control value measurement via the model. That information will be documented separately. In addition, this document describes version 1.0 of the model. FAIR-CAM™ will undoubtedly evolve as we learn more from exercising it, and from additional research.

## Licensing and Use

The FAIR-CAM™ ontology is intended to serve as an international standard for controls physiology. In order to support this objective this work will be licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License (the link can be found at https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode). To further clarify the Creative Commons license related to FAIR-CAM™ content, you are authorized to copy and redistribute the content in its entirety , for non-commercial purposes only, provided that (i) appropriate credit is given to the FAIR Institute, and (ii) a link to the Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License is provided with any distribution of the content. Additionally, if you remix, transform, create derivative works of, or otherwise change and/or build upon the FAIR-CAM™ ontology, you may not distribute the modified materials or use them for commercial purposes. Users of FAIR-CAM™ are required to refer to (http://www.fairinstitute.org/FAIR-CAM/) when referring to the model in order to ensure they are employing the most up-to-date guidance. Commercial use of FAIR-CAM™ is prohibited without the express prior written approval of the FAIR Institute. You are permitted to use the trademark "FAIR-CAM" only as part of a reproduction of the content in compliance with the Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Public License. Please email info@fairinstitute.org with any questions or inquiries.

## A High Level Overview of FAIR

This section provides a high-level overview of FAIR for those readers who are not already familiar with it. Readers who are familiar with FAIR can safely skip this section.

FAIR, or Factor Analysis of Information Risk, is sometimes referred to as a framework, a taxonomy, or an ontology. More simply stated, it is a model of the factors that drive the frequency and magnitude of loss (i.e., "risk") from a loss event scenario. In fact, the FAIR definition for risk is:

*"The probable frequency and probable magnitude of future loss."*

The diagram below illustrates the FAIR model factors and their relationships:

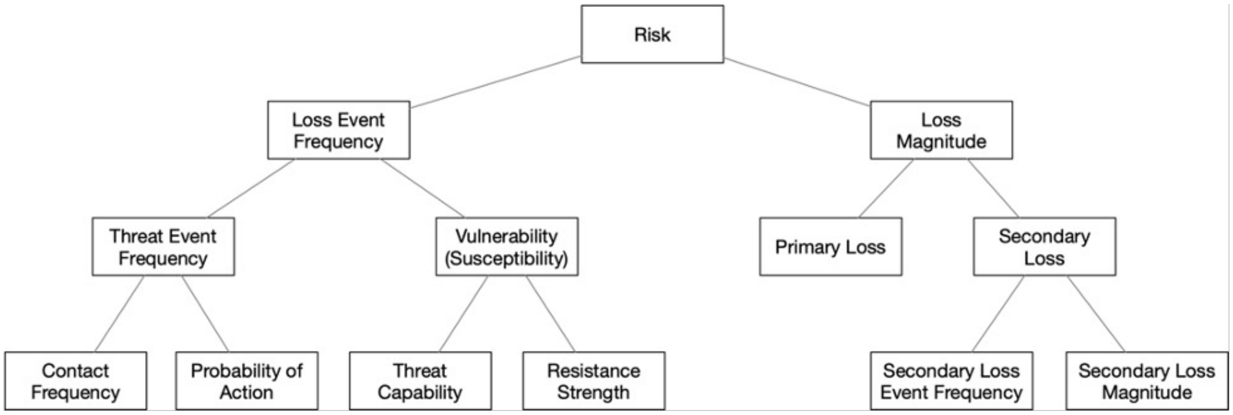

Each node of the model has a very specific definition (provided in other documents available through the FAIR Institute). These explicit definitions and the structure of the model provide many benefits, including:

- Providing a framework for decomposing and simplifying complex risk scenarios
- Reducing reliance on analyst mental models of risk
- Clarifying risk-related discussions (during analysis and when reporting to stakeholders)
- Harmonizing risk analysis between two or more analysts
- Enabling the use of quantification methods (while still supporting qualitative risk measurement)

Also part of the model are definitions for the different forms that loss can take. These definitions simplify the process of gathering, categorizing and reporting loss data:

- Productivity loss
- Response costs
- Replacement costs
- Fines and judgments
- Loss of competitive advantage
- Reputation damage

In order to apply FAIR, you begin by defining a loss event scenario that you want to measure the frequency and magnitude of — e.g., "An outage of a specific business process due to a cybercriminal ransomware attack."

Once the scenario is clearly defined, you can begin to gather data and calibrated estimates for the various FAIR factors and loss forms associated with the scenario. If you're doing a quantitative analysis, your data should be represented as ranges or distributions in order to faithfully represent the uncertainty in your values.

With your input values in hand, you apply them to the model and derive your estimated risk for that scenario. The results can be represented in various ways, both quantitative and qualitative, to best meet decision-maker needs.

FAIR analysis results may be used in various ways, including:

- To prioritize the many risks an organization faces
- To perform cost-benefit analysis in support of proposed risk mitigation improvements
- To defend risk management program elements in the face of budget cuts
- To communicate risk to business stakeholders who are used to quantitative metrics in economic terms

Additional resources regarding FAIR may be found on the FAIR Institute website (www.fairinstitute.org), as well as from the Open Group (www.opengroup.org), which offers a professional certification in FAIR.

## How FAIR-CAM Relates to FAIR

You may have noticed that controls (e.g., passwords, policies, redundant systems, warning signage, etc.) aren't explicitly defined within the FAIR model. This isn't an oversight, but rather reflects the fact that the FAIR model is agnostic and can be applied to any form of risk (e.g., cybersecurity, health and safety, natural disasters, armed conflict, fraud, etc.), while many controls tend to be relevant only to specific types of loss event scenarios.

To-date, personnel performing FAIR analyses have been expected to understand or figure out which controls are relevant to the scenarios they're analyzing, and appropriately account for the effect of those controls on risk. Unfortunately, this can be very challenging without a clear understanding of controls physiology. Furthermore, the absence of a controls physiology model has made it impossible to reliably leverage security telemetry to automate risk measurements.

The FAIR-CAM™ ontology can be thought of as an overlay to FAIR that describes how various control functions affect risk. This will simplify the analysis process, and improve analysis reliability. As will be described later, it also provides a way to accurately associate the elements within common control frameworks with their effect on risk.

# Laying a Basic Foundation

Much of the FAIR-CAM™ model should come across as logical, relatively straight-forward, and well-aligned with your intuitions. Some of it may even feel like "old news." Other parts, however, are nuanced and unobvious, and still others might be perceived as contrary to conventional wisdom or your intuition. This section will focus on conveying some of the basic principles underlying FAIR-CAM™. For now we'll avoid some of the more challenging aspects of controls physiology. Those will come later.

In order to make this first exposure to controls physiology principles as relatable as possible, this section will use an analogy that has nothing to do with cybersecurity or formal risk management. A cybersecurity example will be provided later in the document.

## *A new bicycle*

Imagine for a moment that you are the parent of a six year-old daughter who just received her first "real" bicycle as a birthday present. As a life-long cyclist, you're excited about the prospect of future rides with her through the countryside, but at the same time you recognize there is risk associated with cycling. So now you're faced with the prospect of trying to protect the most important person in your world from this newly added exposure to harm.

As you consider your options for reducing bicycle riding risk to your daughter, you put together a list of controls that you could apply:

- Buying her a bicycle helmet
- Buying one of those brightly colored flags on a stick that attach to the rear of the bicycle
- Putting her through a class on bicycle safety
- Setting restrictions on where/when she can ride
- Installing training wheels on the bicycle (even though she knows how to ride without them)
- Having her wear a brightly colored vest
- Having her wear elbow and knee guards
- Trading the bicycle for an Xbox

Some of these seem like obvious choices, while others are more questionable for one reason or another. Regardless, you're now faced with some additional questions:

- Are some controls more valuable than others?
- How and where do you draw the line?
- Are there other controls you should consider?

Answering the first question requires being able to measure the value of individual controls, while answering the second question requires being able to understand the aggregate effect of two or more controls. As for the third question — the answer is "yes", as we'll see shortly.

## *Defining "value"*

It might seem rhetorical to define "control value", but to minimize the opportunity for confusion here's the definition as used throughout this document:

*"How much risk a control reduces."*

And, if we leverage the FAIR definition for risk — "The probable frequency and probable magnitude of future loss" — we can infer that controls provide value by reducing either the frequency or magnitude of loss events.

With this in mind, the next question we have to consider is, "What are the loss event scenarios each of the controls are relevant to."

There are a surprising number of potential loss event scenarios associated with bicycle riding, but for this exercise we'll focus on just one — your daughter being struck and injured by a motor vehicle while riding her bicycle. This will provide the context for evaluating the risk reduction value of our various control options.

**Key takeaway:**

**Loss event scenarios provide the context for evaluating the risk reduction value of controls.**

## *How controls affect risk*

We mentioned above that controls provide value by affecting the frequency or magnitude of loss, so let's look at our control options in that light:

| Control | Affects Frequency | Affects Magnitude |
|---|---|---|
| **Helmet** | | X |
| **Flag** | X | |
| **Safety class** | X* | |
| **Riding restrictions** | X* | |
| **Training wheels** | X | |
| **Colored vest** | X | |
| **Elbow and knee guards** | | X |
| **Xbox** | X | |

There are, however, a few things to consider about this list:

- Training wheels may be helpful in preventing accidents where your daughter tips over on her bicycle, but they are a lot less relevant to accidents involving motor vehicles.
- Similarly, elbow and knee pads might mitigate scrapes and bruises from a fall, but they aren't probably all that effective in mitigating the severity of accidents involving motor vehicles.
- Although helmets can significantly reduce the odds of serious head and brain injury, they provide no protection to the rest of the rider's body.

**Nuance warning:** You may have noted that "Riding Restrictions" and "Safety class" have an asterisk next to them. That's because neither of these controls directly affect either the probability or magnitude of loss. Instead, they indirectly affect risk by affecting the efficacy of another control — in this case, your daughter. You see, while riding her bicycle, your daughter will make decisions that directly affect risk, like how often she chooses to ride on busy streets, rides during rush hour, etc.[3] In this role, she is the control that directly affects the probability of an accident.

[3] In FAIR terms, your daughter's decision-making regarding when and where she rides her bicycle affects Contact Frequency — i.e., how often she comes into potentially dangerous contact with threat agents.

**Key takeaways:**

- **Controls can affect risk in different ways (e.g., affecting frequency vs. magnitude of loss), as well as by affecting risk directly or indirectly.**
- **Not all controls are relevant to all of the potential loss event scenarios.**
- **Within the FAIR-CAMTM ontology, controls that *directly* affect the frequency or magnitude of loss are referred to as Loss Event Controls (LECs).**

The restrictions you define (should) affect her choices of where and when she rides her bicycle — i.e., the restrictions are a control that *indirectly affects risk* by affecting her performance as a control.

## *Control Performance*

A lot of data exists regarding the efficacy of certified bicycle helmets in preventing serious head trauma, and death. For the sake of simplicity in this exercise, we'll say that wearing a helmet reduces the odds of serious head trauma by 60%. We'll refer to this as the helmet's *"Intended Performance"*.

There are factors, however, that affect how well the helmet actually performs as a control. For example, a helmet that isn't properly fitted or isn't worn properly isn't likely to provide the same level of protection as one that is. And a helmet that sits in a dusty corner of the garage while your daughter is riding her bicycle will reduce the odds of serious head trauma by 0%.

The point is, a control's "Operational Performance" can be less than its Intended Performance. In FAIR-CAM™, when a control is operating in a sub-optimal condition, or isn't operating at all, it's referred to as being in a "Variant Condition."

What we typically want is for Operational Performance to equal Intended Performance as closely as possible. We accomplish this by minimizing the frequency and duration of variant conditions. As we'll see in the next section, this is where we start adding controls to our list.

**Key takeaways:**

- **A control's Operational Performance can be less than its Intended Performance.**
- **Operational Performance is a function of the frequency and duration of variance from a control's intended condition.**

## *Managing Operational Performance*

We've determined that a helmet can reduce the odds of serious head trauma by as much as 60%, but that its actual efficacy in performing this function depends upon how much of the time your daughter is wearing it properly while riding her bicycle. Consequently, you begin making a list of things you can do to affect this:

| Control | Affects the frequency of variance | Affects the duration of variance |
|---|---|---|
| Define clear expectations about wearing the helmet | X | |
| Educate her on those expectations | X | |
| Let her pick out a certified helmet she likes and that fits well | X | |
| Make her do more chores if she's caught riding without the helmet | X | |
| Check up on her periodically while she's riding her bicycle | | X |
| Make her put her helmet on when she's caught riding without it | | X |

There are a few things to consider about the elements within this list:

- The first two controls have a dependency on one another. You may have decided that she must wear the helmet every time she rides her bicycle, but if you don't inform her of that expectation (as well as the consequences for not complying), then the efficacy of having defined expectations is significantly reduced (i.e., variance will be more frequent). Likewise, if you educate her on the fact that the helmet is really important, but your expectations regarding its use (and consequences for non-use) are ambiguous, then variance is likely to be more frequent.
- It's easy to imagine that letting her pick out a helmet that she likes will increase the likelihood that she wears it (i.e., reduce variance frequency). That said, this wouldn't typically be thought of as a "control." In FAIR-CAM™, however, a control is considered to be anything that can be used to reduce the frequency or magnitude of loss (either directly or indirectly). This more inclusive definition enables us to identify and leverage important control opportunities, like allowing her a choice of helmet.[4]

[4] Within a cybersecurity context, working collaboratively with your business stakeholders to identify controls that achieve risk management objectives and yet are a good "fit" tends to be an underutilized opportunity to reduce risk by minimizing variance and improving control Operational Performance.

- Checking-up on your daughter periodically can be considered analogous to auditing (i.e., identifying variant conditions), and making her put her helmet on when found out of compliance is analogous to remediation (i.e., correcting variant conditions).

It's useful to note that some controls require less variance management than others. For example, your daughter has to make a decision every time she gets on her bicycle — does she wear the helmet, or not. On the other hand, that bicycle flag you bolted onto the back of her bicycle doesn't require a decision on her part and is therefore a more reliable control. In other words, its Operational Performance more reliably equals its Intended Performance.

**Key takeaways:**

- **A "Control" is anything that can be used to directly or indirectly affect the frequency or magnitude of loss.**
- **The Operational Performance of controls is strongly affected by variance management controls. This demonstrates one of the ways in which control dependencies are so important to effective controls assessment and risk management.**
- **There are different ways to affect variance — i.e., you can affect its frequency or its duration.**
- **Some variance management controls are dependent upon other variance management controls.**
- **Some controls have an inherently higher level of Operational Performance than others because they require less variance management.**

## *When There's More than One Control*

Here's a question for you to think about. Does the helmet reduce the same amount of risk when it's the only Loss Event Control as it does when combined with another Loss Event Control[5] — for example, if your daughter wears a helmet AND has one of those brightly colored flags attached to her bicycle.

In fact, if the only Loss Event Control your daughter uses is a bicycle helmet, then the helmet is going to reduce more risk than it does when combined with another Loss Event Control. In other words, when we add another control to the mix, the helmet's "risk reduction value" goes down. This may seem counter-intuitive, so let me explain.

Let's imagine that without the bicycle flag there's a 10% probability of your daughter being struck by a car within the next year while riding her bicycle. And if she's properly wearing a helmet, the probability of a severe head injury is reduced by 60%. Purely for the sake of illustration, if we assume that the lifetime cost of treatment for a severe head injury is $1M, then the helmet has lowered your one-year risk from $100,000[6] (inherent risk[7]) to $40,000[8] (residual risk) — i.e., the helmet is reducing risk by $60,000.

[5] Remember that Loss Event Controls (LECs) are those controls that directly affect the frequency or magnitude of loss.

[6] Inherent risk: 10% x $1M = $100,000

[7] For the purposes of this document "inherent risk" means "risk level before a control is applied" rather than the hypothetical "no controls whatsoever risk level".

[8] Residual risk: 10% x ($1M x 40%) = $40,000

What happens, however, if we add a brightly colored flag to her bicycle that cuts in half the probability of her being struck by a car? Well, reducing the probability of an accident by 50% means that the inherent risk (i.e., pre-helmet loss exposure) is $50,000 rather than $100,000 (i.e., 5% probability of an accident x $1M = $50,000). Consequently, the helmet's 60% reduction of serious head injury takes our residual risk to $20,000 (i.e., 40% of $50,000), which means the helmet's risk reduction value is now $30,000 rather than $60,000.

The point here is that in order to understand a control's risk reduction value, we also have to take into account other Loss Event Controls that are in place and relevant to the same loss event scenario(s). The more LECs that are in place for a given scenario, the lower the risk reduction value for each of those controls — i.e., the "risk reduction pie" is being shared by more controls.

**Nuance warnings:**

- Some controls are dependent upon other controls in order to provide any risk reduction value. For example, the remediation control "Make her put her helmet on when she's caught riding without it" isn't going to improve the helmet's Operational Performance if the "Check up on her periodically while she's riding her bicycle" identification control isn't applied as well.
- Although adding more LECs can diminish the risk reduction value of other LEC's for a given loss event scenario (unless they're dependent upon one another, as mentioned above), ***adding*** controls that reduce variance frequency or duration (e.g., clear expectations, auditing, etc.) ***increases*** an LEC's risk reduction value by improving its Operational Performance.

**Key takeaways:**

- **The value of a control depends upon how significantly it affects the frequency or magnitude of loss from one or more loss event scenarios.**
- **The value of a Loss Event Control *decreases* as other Loss Event Controls are implemented that are relevant to the same loss event scenario(s) — unless two or more of those controls are dependent upon one another.**
- **The value of a Loss Event Control *increases* as controls are added that reduce variance frequency or duration — i.e., as its Operational Performance improves.**

- However, just as adding more LECs can reduce the risk reduction value of other LEC's, there can be diminishing returns to improving Operational Performance from piling on more variance management controls.

## *A brief review of decision-making*

Earlier, we discussed how defining restrictions for when and where your daughter could ride her bicycle, and educating her on those restrictions, indirectly affects risk by influencing her decision-making. In that instance the operator of the bicycle, your daughter, is the Loss Event Control. A similar example within cybersecurity is when personnel make decisions regarding whether or not to open email attachments. There, too, the person is acting as a control, and the ***quality of their decisions determines their efficacy as a control.***

We also discussed the fact that every time your daughter rides her bicycle, she has to make an explicit decision to wear, or not wear, her helmet.  Those decisions affect the Operational Performance (and thus, the value) of the helmet as a Loss Event Control.

There are, however, many other decisions taking place that affect risk within this scenario.  A partial list includes:

- You're making decisions about what the riding restrictions should be.
- You're making decisions about what the consequences will be when she doesn't wear her helmet, or when she violates the riding restrictions.
- You're making decisions about which other Loss Event Controls (e.g., the flag) to employ.
- You're making decisions about how often to check up on her to see if she's wearing her helmet, and to confirm that she's only riding when/where she should be.

We'll get into a lot more detail elsewhere regarding the factors that affect decision-making.  For now, simply recognize that decision-makers can affect risk in different ways.  Sometimes the decision-makers, themselves, are acting as Loss Event Controls.  Other times, decision-makers are affecting the Operational Performance of controls (like a choice to wear the helmet).  And still other times, decisions are being made that influence other decisions (e.g., defining consequences for non-compliance).

**Key takeaways:**

- **Decisions can directly affect risk when the decision-maker is acting as a Loss Event Control.**
- **Decisions can indirectly affect risk by affecting the Operational Performance of controls.**
- **Decisions can indirectly affect risk by affecting other decisions.**

Fortunately, regardless of the decision-making context, the factors that affect decision-making quality are constant. What these factors are will be covered in an upcoming section on Decision Support Controls.

## *Wrapping up the bicycle scenario*

Of course, parents all over the world have wrestled the bicycle riding scenario more or less successfully for generations, and have done so without something like FAIR-CAM™.  But managing the risk associated with cybersecurity is vastly more complicated than our bicycle riding scenario.

It shouldn't be too difficult to extrapolate from what was discussed in this scenario, to the scores of cybersecurity Loss Event Controls organizations use to manage the frequency and magnitude of the cybersecurity loss event scenarios they face.  Add to that, the many controls required to manage the Operational Performance of controls, as well as the myriad decisions being made — explicitly and implicitly — about organization's' cybersecurity risk landscape.  It should be pretty clear that common control frameworks and mental models that might suffice for something as simple as the bicycle scenario can't be relied upon for cybersecurity.

# Model Overview

This section will describe the formal structure of the FAIR-CAM™ ontology, introducing the model elements and how they fit together. This, too, is intended to be relatively high-level, and will not include details like units of measurement, in-depth guidance regarding how to apply the model, etc. Those details can be found in separate documents, specifically — A Description of the FAIR Controls Analytics Model and Applying the FAIR Controls Analytics Model.

## Three foundational terms

Because ambiguous terminology can complicate almost any discussion, it's important to make certain that we share a common understanding of several key terms as they're used in the FAIR-CAM™ ontology.

### *Control*

*"Anything that can be used to reduce the frequency or magnitude of loss."*

Controls can be: laws, regulations, policies, standards, processes, technologies, people, software, physical structures, etc.. This definition is intentionally broad in scope, as it enables us to account for the risk reduction effects of more things.

### *Control Function*

*"How a control directly or indirectly affects the frequency or magnitude of loss."*

A few examples of how controls can affect risk include:

- Limiting contact with threats (threat avoidance)
- Making it more difficult for threats to adversely affect assets (loss event resistance)
- Providing evidence that a loss event has occurred (loss event visibility)
- Restoring operations after an outage-related loss event has occurred (loss event resilience)
- Reducing the frequency of missing or deficient controls (variance prevention)
- Detecting that controls are missing or in a deficient condition (variance identification)
- Remediation of deficient controls (variance correction)

### *Functional domains*

*"High-level control function categories"*

Functional domain categories distinguish between control functions that affect risk directly, versus those that affect the Operational Performance of controls, versus those that affect decision-making. Two of FAIR-CAM's functional domains were introduced during the bicycle scenario — Loss Event Controls (LECs), and Variance Management Controls (VMCs). We also mentioned the existence of a third one — Decision Support Controls (DSCs). As the name implies, DSCs affect decision-making quality.

The diagram below illustrates the relationships between these functional domains, as well as their relationship with risk:

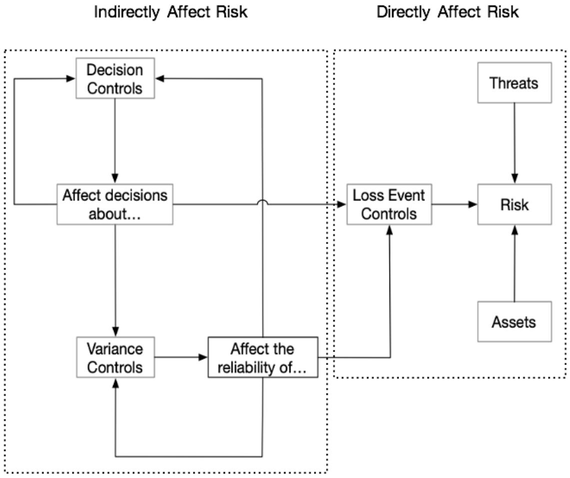

This simple diagram does a good job of showing how the interaction of variance controls and decision controls can create complexity, as well as how they tie back to loss event controls and risk.  This complexity will become even more evident as we get deeper into the model.

Fortunately, by explicitly defining the relationships between these different functional domains and the functions themselves, we have the keys to understand and more effectively manage the controls we apply to our risk landscape.

The rest of this section provides three descriptive resources for each of the functional domains:

1. A diagram that illustrates the control functions and where they fit within the functional domain
2. A table that provides a brief description of each control function in the domain.
3. A flow diagram that illustrates how these functions work to affect risk.

Note that more detailed descriptions of the functional domains and control functions will be provided in a separate document.

## Loss Event Control Functions

As implied by its name, LEC functions are realized by applying controls that directly affect the frequency or magnitude of loss events.

The diagram below illustrates the functional elements within this part of FAIR-CAM™, and their relationships:

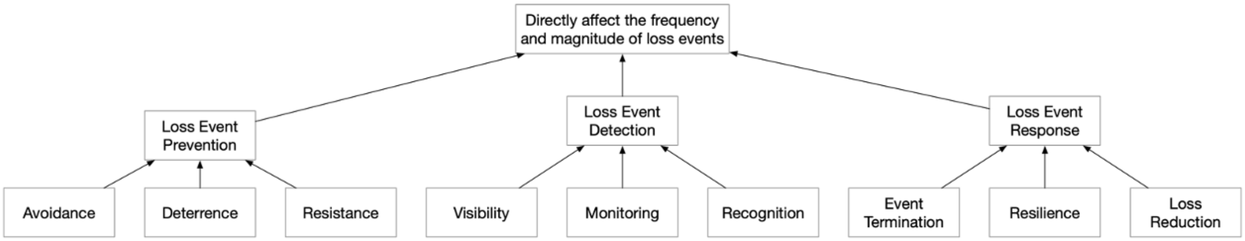

The table below provides a brief description for each of these functions.

| Control Function | Function Description |
|---|---|
| **Avoidance** | Reduce the frequency of contact between threat agents and the assets they might adversely affect. |
| **Deterrence** | Reduce the probability of malicious or unauthorized actions after a threat agent has come into contact with an asset. |
| **Resistance** | Reduce the likelihood that a threat agent's act will result in a loss event. |
| **Visibility** | Provide evidence of activity that may be anomalous or illicit in nature. |
| **Monitoring** | Review data provided by Visibility controls. |
| **Recognition** | Enable differentiation of normal activity from abnormal activity. |
| **Event Termination** | Enable termination of threat agent access or activities that could continue to be harmful. |
| **Resilience** | Maintain or restore normal operations. |
| **Loss Reduction** | Reduce the amount of realized losses. |

The diagram below illustrates how these control functions affect the frequency or magnitude of loss within a loss event scenario.

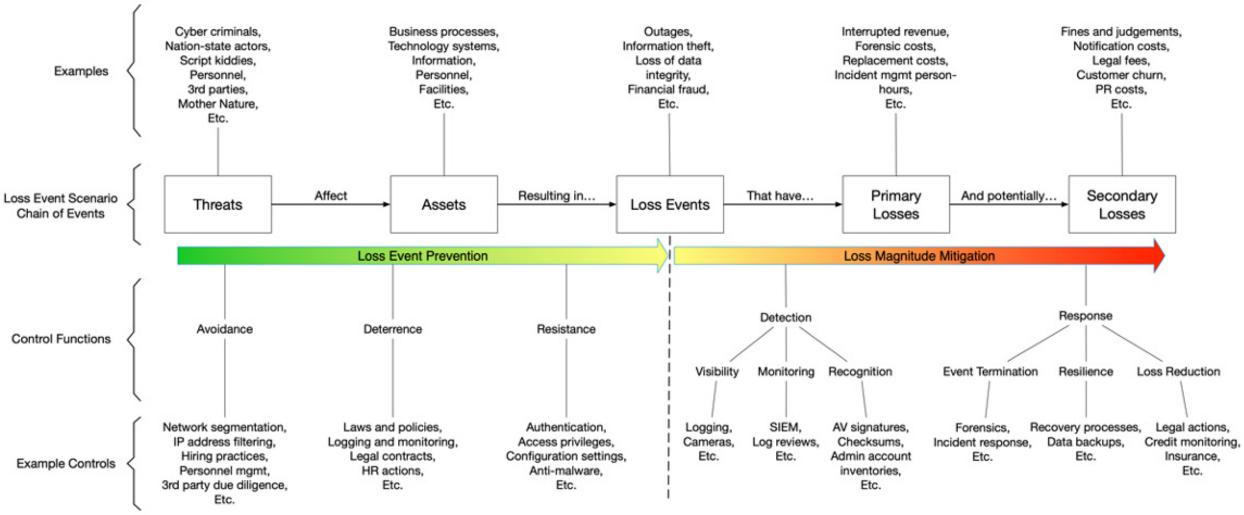
Examples
Cyber criminals, Nation-state actors, Script kiddies, Personnel, 3rd parties, Mother Nature, Etc.
Business processes, Technology systems, Information, Personnel, Facilities, Etc.
Outages, Information theft, Loss of data integrity, Financial fraud, Etc.
Interrupted revenue, Forensic costs, Replacement costs, Incident mgmt person-hours, Etc.
Fines and judgements, Notification costs, Legal fees, Customer churn, PR costs, Etc.
Loss Event Scenario Chain of Events
Threats
Affect
Assets
Resulting in...
Loss Events
That have...
Primary Losses
And potentially...
Secondary Losses
Loss Event Prevention
Loss Magnitude Mitigation
Control Functions
Avoidance
Deterrence
Resistance
Detection
Visibility
Monitoring
Recognition
Response
Event Termination
Resilience
Loss Reduction
Example Controls
Network segmentation, IP address filtering, Hiring practices, Personnel mgmt, 3rd party due diligence, Etc.
Laws and policies, Logging and monitoring, Legal contracts, HR actions, Etc.
Authentication, Access privileges, Configuration settings, Anti-malware, Etc.
Logging, Cameras, Etc.
SIEM, Log reviews, Etc.
AV signatures, Checksums, Admin account inventories, Etc.
Forensics, Incident response, Etc.
Recovery processes, Data backups, Etc.
Legal actions, Credit monitoring, Insurance, Etc.

## Variance Management Control Functions

The Intended Performance of controls is the maximum, or ideal, ability of a control to perform one or more control functions. However, a control's Operational Performance can have a huge effect on a control's actual risk reduction value. In fact, in many cases a control with a lower Intended Performance but greater Operational Performance will reduce more risk than a control with higher Intended Performance but lower Operational Performance.  For example, a bicycle helmet that is 100% effective at preventing head injury (hypothetical, of course) but that your daughter refuses to wear because it's uncomfortable and ugly, will reduce less risk than a helmet that's only 60% effective, but is worn consistently because she picked it out.

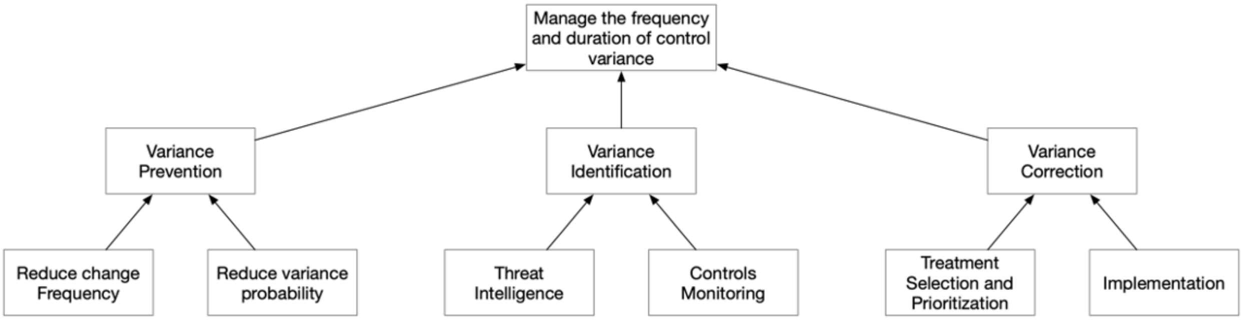

Consequently, even though they affect risk indirectly, variance controls are crucial to managing the frequency and magnitude of loss.

The table below provides a brief description for each of the VMC functions.

| Control Function | Function Description |
|---|---|
| **Reduce change frequency** | Reduce the frequency of changes that might introduce variant control conditions. |
| **Reduce variance probability** | Reduce the probability of variance being introduced when changes to systems, networks, etc. occur. |
| **Threat intelligence** | Enable the recognition of changes in the threat landscape that result in loss event controls no longer being as effective as intended. |
| **Controls monitoring** | Enable the recognition that variant control conditions exist. |
| **Treatment and selection prioritization** | Ensure that effective remediation activities are appropriately prioritized. |
| **Implementation** | Correct variant control conditions. |

The diagram below illustrates how these control functions affect the frequency or duration of variant conditions.

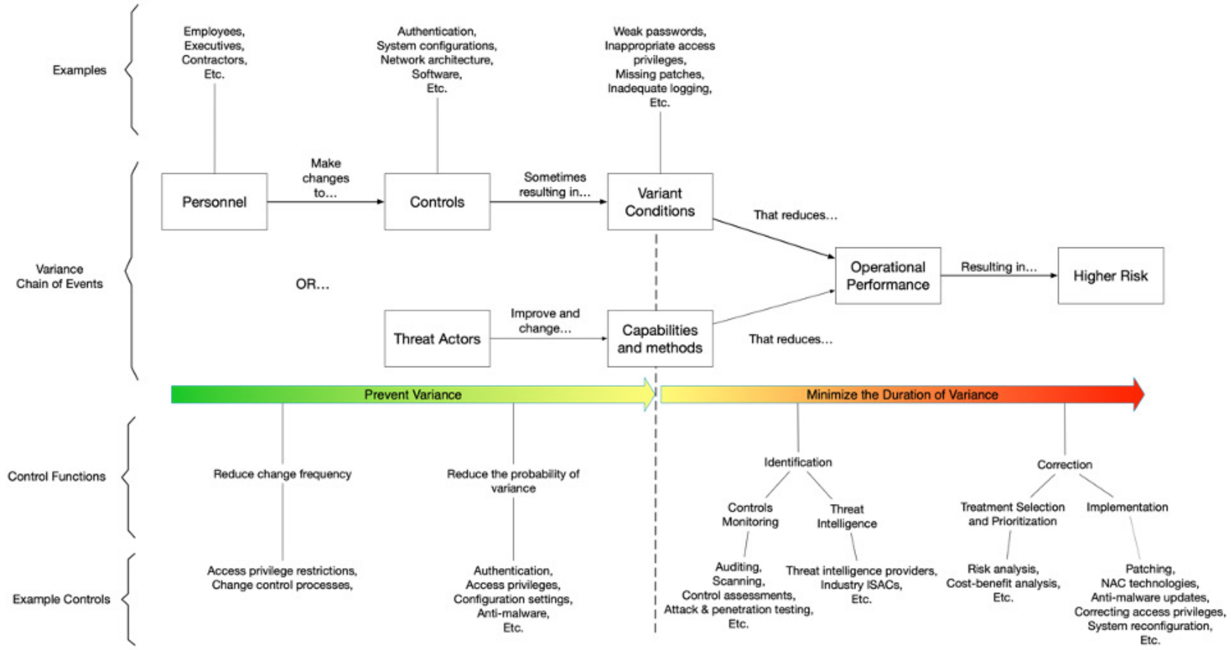

## Decision Support Control Functions

It's probably obvious by now that the effects of decision-making can range from being very local and tactical (e.g., open that email attachment, or not) to global and strategic (e.g., setting the organization's risk appetite). Furthermore, there are multiple factors that can affect how closely aligned decisions are with the organization's expectations and objectives. Consequently, the Decision Support functional domain is the most complicated and far-reaching part of the controls physiology landscape and FAIR-CAM™.

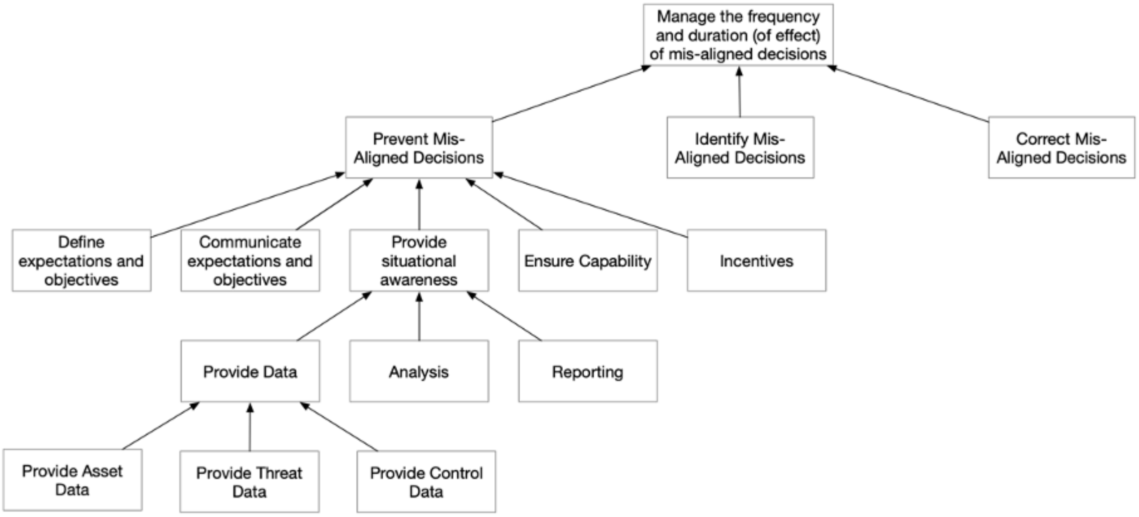

The table below provides a brief description for each of the DSC functions.

| Control Function | Function Description |
|---|---|
| **Define expectations and objectives** | Define an organization's risk management expectations and objectives. |
| **Communicate expectations and objectives** | Ensure that responsible persons are aware of and understand the organization's risk management objectives and priorities. |
| **Asset data** | Provide data related to the assets at risk. |
| **Threat data** | Provide data related to threats against assets. |
| **Control data** | Provide data related to controls. |
| **Analysis** | Accurately synthesize asset, threat, and control data for decision-makers. |
| **Reporting** | Provide analysis results to decision-makers in a manner that meets their needs. |
| **Provide Capability** | Ensure that the decision-maker has the necessary skills, authority, and resources to make decisions that are aligned with the organization's expectations and objectives. |
| **Incentives** | Ensure that personnel are motivated on a personal level to make decisions that are aligned with the organization's expectations and objectives. |
| **Mis-aligned decision identification** | Identify decisions that are not aligned with organization expectations or objectives. |
| **Mis-aligned decision correction** | Correct decisions that were not aligned with organization expectations or objective. |

The diagram below illustrates how these control functions affect the frequency or duration of effect of mis-aligned decisions.

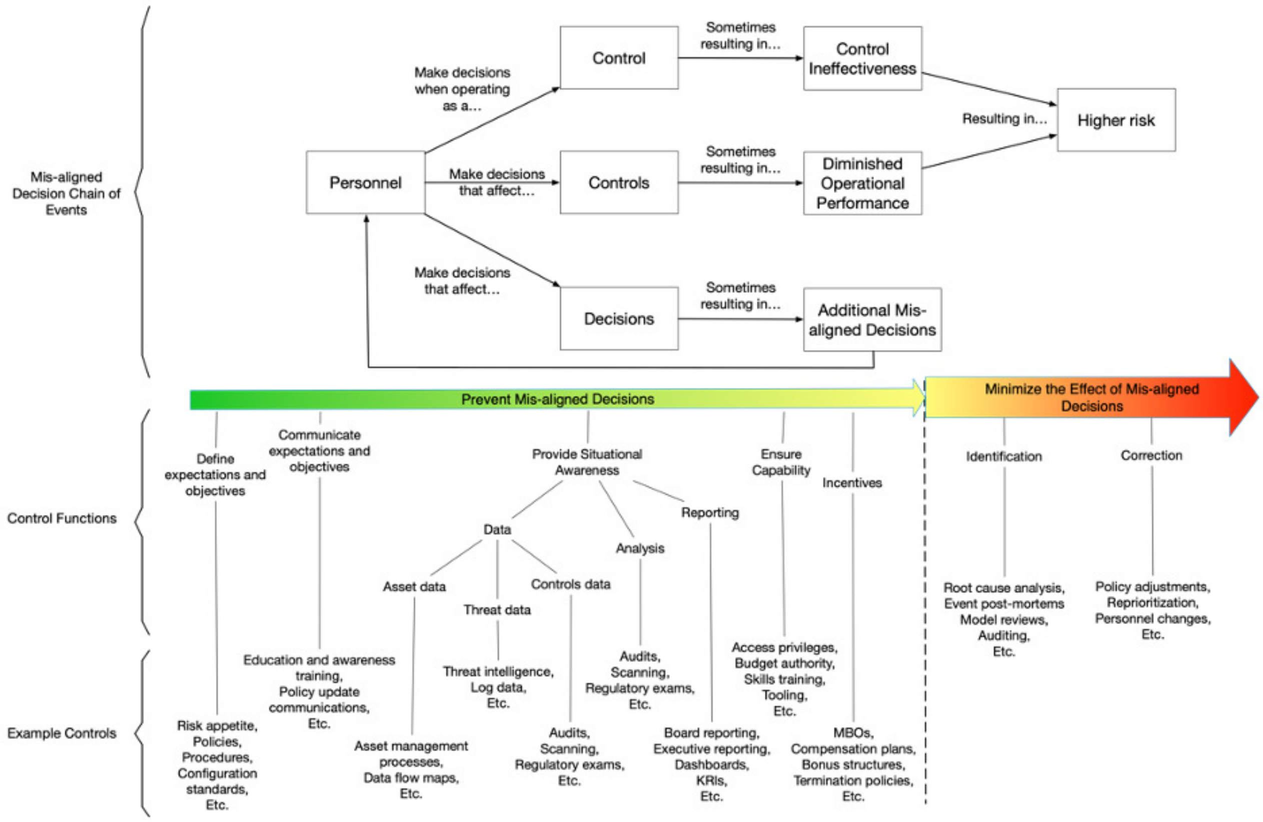

Mis-aligned Decision Chain of Events
Personnel
Make decisions when operating as a…
Control
Sometimes resulting in…
Control Ineffectiveness
Make decisions that affect…
Controls
Sometimes resulting in…
Diminished Operational Performance
Resulting in…
Higher risk
Make decisions that affect…
Decisions
Sometimes resulting in…
Additional Mis-aligned Decisions
Prevent Mis-aligned Decisions
Minimize the Effect of Mis-aligned Decisions
Control Functions
Define expectations and objectives
Communicate expectations and objectives
Provide Situational Awareness
Data
Asset data
Threat data
Controls data
Analysis
Reporting
Ensure Capability
Incentives
Identification
Correction
Example Controls
Risk appetite, Policies, Procedures, Configuration standards, Etc.
Education and awareness training, Policy update communications, Etc.
Asset management processes, Data flow maps, Etc.
Threat intelligence, Log data, Etc.
Audits, Scanning, Regulatory exams, Etc.
Audits, Scanning, Regulatory exams, Etc.
Board reporting, Executive reporting, Dashboards, KRIs, Etc.
Access privileges, Budget authority, Skills training, Tooling, Etc.
MBOs, Compensation plans, Bonus structures, Termination policies, Etc.
Root cause analysis, Event post-mortems Model reviews, Auditing, Etc.
Policy adjustments, Reprioritization, Personnel changes, Etc.

# A Cybersecurity Example

Understanding the risk reduction value of controls is fundamental in order to reliably answer many risk management questions, for example:

- Are we sufficiently protected against a specific loss event scenario?
- Which control deficiency or gap is most important to fix?
- Which of two or more risk mitigation measures provides the most value?
- Can we retire one or more controls without significantly increasing our risk?
- Have we reduced risk enough overall?
- Have we reduced too much risk overall?

We'll use the first of these question use-cases for our example. Note that because this is meant to be an introduction, we'll avoid the use of numbers and math.

## *The question: Are we sufficiently protected against ransomware?*

Recall that loss event scenarios provide the context for measuring control risk reduction value. Furthermore, because clarity is crucial when it comes to measurement, we need to be explicit about the scope of what we're analyzing. With that in mind, the following is a description of our scenario:

- The assets that we most care about are systems that support key business functions.
- The threat community is cybercriminals.
- The outcome that we want to avoid is an outage of key business processes.
- The most likely initial point of attack is one or more user end-points (laptops, desktops, etc.).
- The most likely initial method of attack will be phishing.

We could define a different ransomware scope for our analysis by making different assumptions (e.g., choosing a different initial point of attack), but these will suffice for this example.

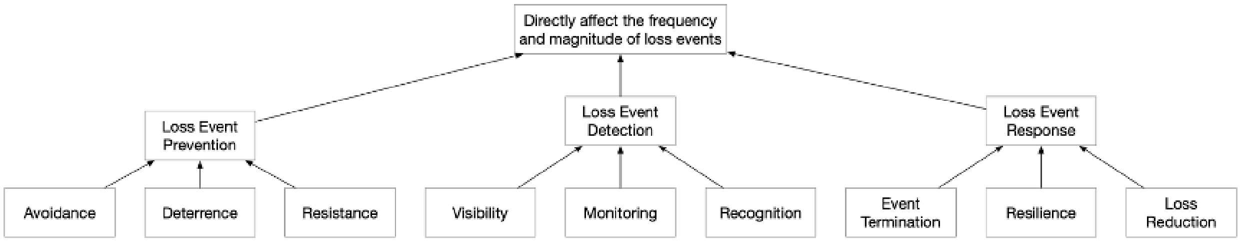

## *A look at Loss Event Controls*

The table below provides a list of Loss Event Controls that are relevant to this loss event scenario and assigns them to the Loss Event Control functions they serve. Note that your organization may have fewer, more, or different controls than are shown in this example. These were chosen simply for illustration purposes.

| Functional Category | Control Function | Controls |
|---|---|---|
| **Loss Event Prevention** | **Avoidance** | • Email anti-malware technology<br>• URL filtering |
| | **Deterrence** | N/A |
| | **Resistance** | • The end-point users themselves<br>• An anti-malware technology<br>• The end-points themselves (their patching levels and configuration)<br>• The patching level and configuration of software on the end-points (e.g., email, browser, etc.)<br>• User privilege restrictions |
| **Loss Event Detection** | **Visibility** | • Anti-malware<br>• Host IDS |
| | **Monitoring** | • Anti-malware<br>• Host IDS<br>• SIEM |
| | **Recognition** | • Anti-malware (signatures)<br>• Host IDS (signatures and heuristics)<br>• SIEM (data analysis) |
| **Loss Event Response** | **Event Termination** | • Anti-malware (sandboxing, etc.)<br>• Forensics<br>• Incident response processes<br>• System segregation<br>• System rebuild |
| | **Resilience** | • Backups<br>• Data recovery technologies and processes |
| | **Loss Reduction** | • Insurance |

A few notes about the table above:

- Avoidance controls minimize contact between the attacker and the point of attack (the end-point).
- Just as your daughter was a Loss Event Control in our bicycle scenario, the end-point user is a Loss Event Control against a phishing attack.
- The end-point operating systems themselves are Loss Event Controls. This may run counter to your intuition, but the explanation will require more time than we can take here. The reasoning will be explained in a separate document. For the moment, please trust that this is necessary and will make sense.
- As with the point above, most people don't think of software as a control. Here again, an explanation will be provided in a separate document.
- You'll note that some cybersecurity solutions perform multiple control functions (e.g., anti-malware). Their performance in these separate functions can be measured distinctly.

Each of these controls will have an Intended Performance level — i.e., their maximum level of efficacy given how an organization has chosen to implement them (configuration standards, etc.). Note that the word "chosen" is underlined in the previous sentence. This is to highlight the fact that these choices represent decisions the organization made — decisions that were supported (well or poorly) by Decision Support Controls, which will be discussed later. You'll see this word underlined elsewhere for the same reasons.

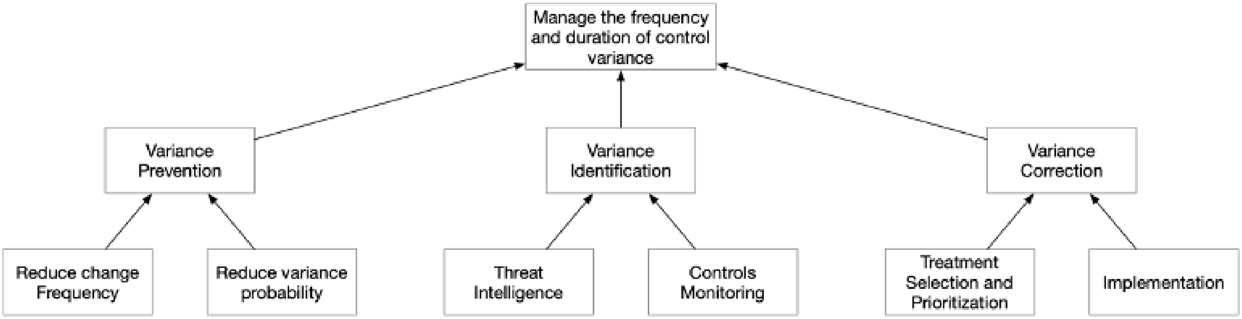

## *A look at Variance Management Controls*

The Operational Performance of our Loss Event Controls will be significantly affected by the Variance Management Controls we use to limit the frequency and duration of variant conditions.

The table below lists some of the VMCs our hypothetical organization has chosen to put in place for the ransomware-relevant LECs:

| Functional Category | Control Function | VM Controls | Affected LECs |
|---|---|---|---|
| **Variance Prevention** | **Reduce Change Frequency** | • User privilege restrictions | • System and software configurations |
| | | • Change control | • System and software configurations<br>• Anti-malware<br>• Host IDS |
| | **Reduce Variance Probability** | • Automation | • System and software configurations |
| | | • Decision support (e.g., phishing training, policy awareness, etc.) | • Personnel |
| **Variance Identification** | **Threat Intelligence** | • Involvement in an industry ISAC | • System and software configurations<br>• Host IDS<br>• Anti-malware<br>• Forensics<br>• Incident response |
| | | • Internal threat monitoring | • System and software configurations<br>• Host IDS<br>• Anti-malware<br>• Forensics<br>• Incident response |
| | **Controls Monitoring** | • Auditing | • All of the LECs |
| | | • Scanning | • Systems and software configurations |
| | | • Attack & penetration testing | • All of the LECs |
| | | • Phishing testing | • Personnel |
| **Variance Correction** | **Treatment Selection and Prioritization** | • Decision support (e.g., risk appetite, policies, risk analysis, incentives, etc.) | • All of the LECs |
| | **Implementation** | • Adjusting access privileges | • Systems and software configurations |
| | | • Reconfiguring systems or software | • Systems and software configurations |
| | | • Patching | • Systems and software configurations |
| | | • Decision support (e.g., remedial phishing training) | • Personnel |
| | | • Implementation of new controls | • Depends on the control gap |

A few notes about the table above:

- Some VMCs affect multiple LECs.
- There may be — intentionally or otherwise — some redundant VMCs (e.g., involvement in ISACs and Internal threat monitoring both provide Threat Intelligence). This redundancy may just be to provide resilience, or because there are subtle differences in how they perform their control function. Or, of course, redundancy may exist simply because nobody had really given it much thought.
- If VMCs apply to more than one LEC, they may be applied differently for one reason or another. For example, how frequently an organization audits one LEC may differ from how frequently they audit another LEC.
- Decision Support Controls come into play in both preventing and correcting variance.

Measuring the risk reduction value of our LECs requires that we understand the frequency and duration of variant conditions. In some cases, we'll have good empirical data about variance from sources like cybersecurity scanning technologies or audit history. In other cases, especially in the early days of applying FAIR-CAM™, we might need to make estimates regarding the expected frequency and duration of variance based upon the Variance Management Controls that are in place. Regardless, recognizing how variance affects control performance it becomes obvious that the frequency and duration of variance should be considered key metrics.

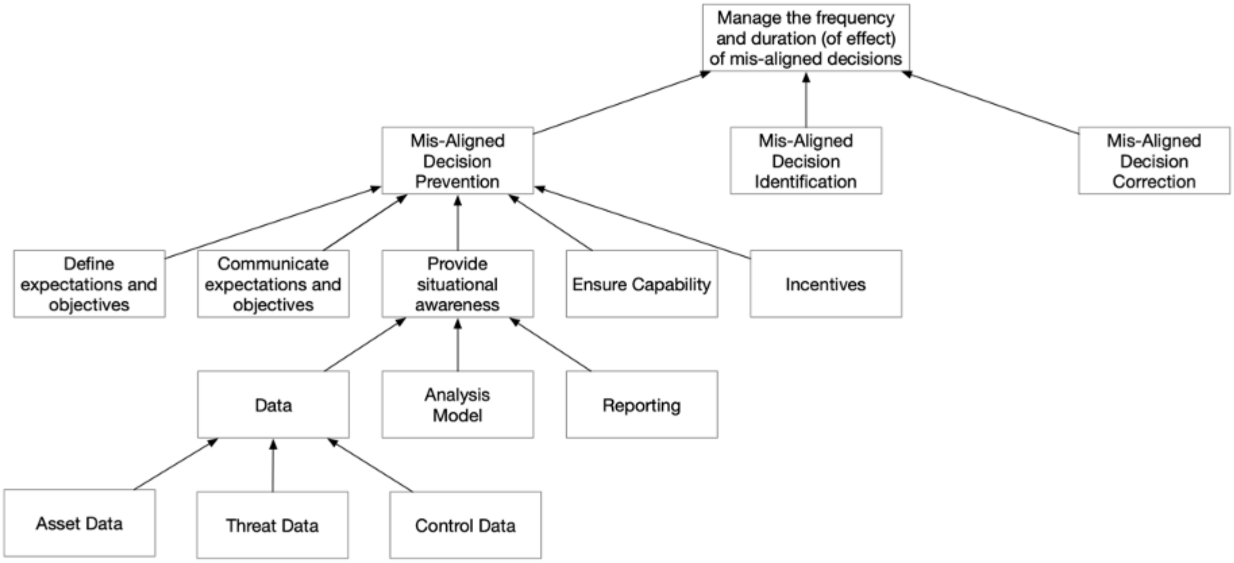

Manage the frequency and duration (of effect) of mis-aligned decisions
Mis-Aligned Decision Prevention
Mis-Aligned Decision Identification
Mis-Aligned Decision Correction
Define expectations and objectives
Communicate expectations and objectives
Provide situational awareness
Ensure Capability
Incentives
Data
Analysis Model
Reporting
Asset Data
Threat Data
Control Data

## *A look at Decision Support Controls*

It's obvious that decisions are central to, and interwoven throughout, any cybersecurity and risk management program. And invariably, the better-aligned those decisions are with the organization's expectations and objectives, the more successful the organization will be.

The table below lists a few of the decisions that surround our example scenario and categorizes them in terms of the role they play in managing risk:

| How they affect risk | Example Decisions |
|---|---|
| **Directly, when a human is acting as a LEC.** | • Personnel deciding whether to click on an email attachment<br>• Personnel deciding whether to override a browser's warning about a potentially malicious website |
| **Indirectly, by affecting the Operational Performance of LECs.** | • Personnel deciding whether to install unauthorized software<br>• Personnel deciding whether to make unapproved configuration changes to their system<br>• Personnel deciding whether to disable an annoying security setting<br>• Personnel choosing their passwords<br>• Management deciding whether/when to patch vulnerable conditions<br>• Choosing which VMCs to apply, and how to apply them |
| **Indirectly, by affecting other decisions.** | • Setting the organization's risk appetite<br>• Defining policies and expectations (e.g., regarding email attachments, unauthorized software, etc.)<br>• Communicating policies and appetite<br>• Defining decision-making authority<br>• Resource allocations |

The first two categories are relatively straightforward in terms of how they can be applied within risk analysis. When a person is acting as a Loss Event Control, their efficacy is a function of the decisions they make. When personnel or management are making decisions that affect the Operational Performance of a LEC, this will affect the frequency or duration of LEC variance.

The third category, however, is a bit less obvious. For example, when decisions are being made regarding the organization's objectives and expectations (e.g., policies), this often establishes the Intended Performance of controls (e.g., setting password requirements). In other cases, however, these decisions will affect the Operational Performance of controls (e.g., setting patch frequency requirements). Either way, as long as you know how the decision is affecting performance, you can appropriately factor it into an analysis.

Furthermore, by understanding the relationships and dependencies between controls we can perform much stronger root cause analyses. This will enable us to reduce or eliminate the "risk management groundhog day" phenomena of having to fight the same control deficiency problems over and over.  By the way, decision support control weaknesses are invariably at the root of risk management groundhog days.

## *Evaluating our defensive posture*

We've defined our scenario and identified the relevant LECs, VMCs, and DSCs.  This by itself may have provided a much clearer sense of our ransomware defenses, but how do we use that information to measure how much risk remains, and more confidently decide whether we need to do more?  Here again, the actual quantitative analysis can get complicated very quickly.  Too complicated for an introductory document like this.  So for now we'll just discuss the steps involved.

The first step is to determine what the Intended Performance is for each LEC — i.e, how effectively it is supposed to perform the function(s) it serves.  We follow that by understanding the frequency and duration of variance for each LEC, which allows us to derive a control's Operational Performance.  These values may be based on empirical testing and measurement, or estimates.  Early in the use of FAIR-CAM™, we're likely to rely more heavily on estimates while we begin to establish the processes required to gather empirical data.

Once we have those values in hand they can be applied to the appropriate point in the FAIR model (e.g., susceptibility — a.k.a. resistance strength).  If there are multiple LECs for a given control function (e.g., end-point users, AV, etc. for Resistance), we aggregate their effect.  Once that's been done for all of the LECs, we can complete the FAIR analysis to understand how much ransomware risk we have.

Based on the results we may choose to deploy additional LECs.  Alternatively, we may add VMCs to make the Operational Performance of our existing LECs more robust.  Or, we may do both.  It's also not impossible to imagine that we might discover controls that can be retired, which would preserve resources for other control improvement opportunities.  Regardless, we'll be in a much stronger position to make these decisions and defend them.

Before we wrap up this section there is one more dimension of the problem that warrants discussion.  Specifically, we've talked about measuring the Intended and Operational Performance of controls, but what about the parts of the risk landscape that you're in some way blind to?  For example, shadow IT and some third parties.  After all, if you don't know about the existence of technology components and software that's in use then how do you factor it into our analysis?  And even if you know of the existence of these assets but have little to no good data regarding their controls, then how do you deal with that?

The best approach we've found to-date for dealing with blind spots is to analyze them separately.  Limit the scope of your primary analysis to the part of the landscape that you have reasonable visibility into.  If desired, you can then do a second analysis where the input values for asset value/liability, threats, and controls have much greater uncertainty, reflected by broader and flatter input distributions.  In this way you're able to account for the fact that your visibility into the risk landscape isn't complete, and at the same time avoid introducing a lot of uncertainty into your primary analysis.  This approach also provides an effective way to compare and communicate the condition of your "known universe" versus the poorly known parts of your risk landscape.

## *Conclusion*

In every instance where a successful and highly impactful ransomware attack has occurred, multiple control deficiencies have existed.  If this weren't the case, then the attacks simply would not have been successful.  By understanding and applying controls physiology principles and the FAIR-CAM™ ontology, organizations will be able to more reliably establish and maintain controls that effectively protect them against ransomware and other cyber, technology, and operational risks.[9]

[9] Even though this document has focused primarily on the cybersecurity context, controls physiology and FAIR-CAM ontologies apply well to any risk management domain.

It's worth noting that once an organization has profiled its controls landscape in terms of Intended Performance and Operational Performance, as well as which controls affect other controls, then the process of performing risk analysis becomes much faster and simpler.  All you have to do is understand which LECs are relevant to the scenarios and questions you're faced with.  Because all (or at least most) of the control relationships and dependencies will have been established by profiling your controls there will be less work to do for subsequent analyses.

# Mapping Other Frameworks

Anatomy versus physiology seems to be an accurate analogy for understanding the relationship between common control frameworks and the FAIR-CAM™ ontology. As such, FAIR-™ is complementary to these frameworks — and in fact fills a critical gap — rather than displaces them. Unfortunately, even though they are complementary, at this time the relationship could be considered somewhat "challenged" in nature.

## *Mapping Challenges*

Ideally, we should be able to directly measure the risk reduction value of each element in a controls framework — for example, 13.7 Manage USB Devices from the CIS Controls. But if we examine 13.7 closely we'll see that the description includes two elements; "system configuration software" (to manage system settings), and "maintaining an inventory of systems with USB capabilities." These two control elements perform entirely different risk reduction functions, having different units of measurement and different relationships with other controls. Consequently, there's no way to simply map CIS 13.7 Manage USB Devices to FAIR-CAM and measure its value. In fact, many elements in today's control frameworks are defined too imprecisely to reliably measure their value as controls.

## *Control outcomes vs. controls*

A similar challenge exists between FAIR-CAM™ and control frameworks (like NIST CSF) which focus on control outcomes rather than the controls themselves. For example, the description for NIST CSF PR.AC-1 includes multiple control outcomes:

- Issuing credentials
- Verifying credentials
- Revoking credentials, and
- Auditing credentials

Each of these is accomplished by various policy, process, and technology controls, serving distinctly different control functions. Conflating them makes it nearly impossible to measure their value reliably.

The bottom line is that to enable reliable measurement of control efficacy and value, control frameworks can't conflate controls that serve different risk reduction functions.

## *Some good news too…*

Some control frameworks are easier to map to the FAIR-CAM™ ontology than others. Generally speaking, the more granular the control framework, the less frequently you'll encounter the conflation problems described above. For example, NIST 800-53's higher granularity increases the percentage of controls that can be cleanly mapped and effectively measured. Even within less granular frameworks though, some of the control elements will map cleanly to FAIR-CAM™. For example, CIS 1.7 Deploy Port Level Access Control maps very cleanly to the FAIR-CAM™ LEC/Avoidance function. In fact, roughly 85% of the CIS controls map reasonably well to FAIR-CAM™.

Despite the challenges, efforts are underway to map the most commonly used control frameworks to FAIR-CAM. These mappings will be published separately as they're completed.

# Wrapping Up and Next Steps

Hopefully, this document has not only provided a high-level understanding of what FAIR-CAM™ is, but also why controls physiology, as a paradigm and discipline, is necessary and provides value. What we've covered here may also have validated some of your own intuitions, provided you with a deeper appreciation of the complex nature of the controls landscape, as well as perhaps lifted some of the fog surrounding how controls affect risk. Of course, for some long-tenured professionals, this may simply appear to be a formalization of how they intuitively think about the controls landscape.

As was mentioned earlier, the complexity described here doesn't come from a desire to make the problem seem complicated. The problem space is inherently complicated. Regardless, it should be clear that controls assessments and controls management practices which fail to account for this complexity will provide unreliable results.

Something else to keep in mind is that controls physiology introduces an almost entirely new discipline to the risk management profession. This should not only improve our efficacy as a profession, but also will provide new opportunities from a technology solution and consulting services perspective. It also could and should eventually influence risk management regulations and standards.

## *Additional resources*

A training program for FAIR-CAM™ will be available beginning in early 2022. For more information and a schedule of classes, please refer to the FAIR Institute website at fairinstitute.org.

In the meantime, the following documents related to the FAIR-CAM™ ontology are, or will soon be, available from the FAIR Institute for those who would like to know more about the model's structure and application:

- Description of the FAIR Controls Analytics Model Standard
- Applying the FAIR Controls Analytics Model
- A Map of CIS Controls to the FAIR Controls Analytics Model
- A Map of NIST CSF to the FAIR Controls Analytics Model
- A Map of Mitre Att&ck D3fend to the FAIR Controls Analytics Model
- A Map of NIST 800-53 to the FAIR Controls Analytics Model
- A Map of HITRUST CSF to the FAIR Controls Analytics Model
- A Map of ISO2700n to the FAIR Controls Analytics Model
- A Map of COBIT to the FAIR Controls Analytics Model
- A Map of PCI DSS to the FAIR Controls Analytics Model