Standards for Risk Management

Table of Contents

National Institute of Standards and Technology Special Publications

CISA | CYBERSECURITY AND INFRASTRUCTURE SECURITY AGENCY

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY SPECIAL PUBLICATIONS

CISA
CYBER+INFRASTRUCTURE

1

**001 Instructor: We're going to talk about what we might find in terms of standards for risk management.

NIST SP 800-30

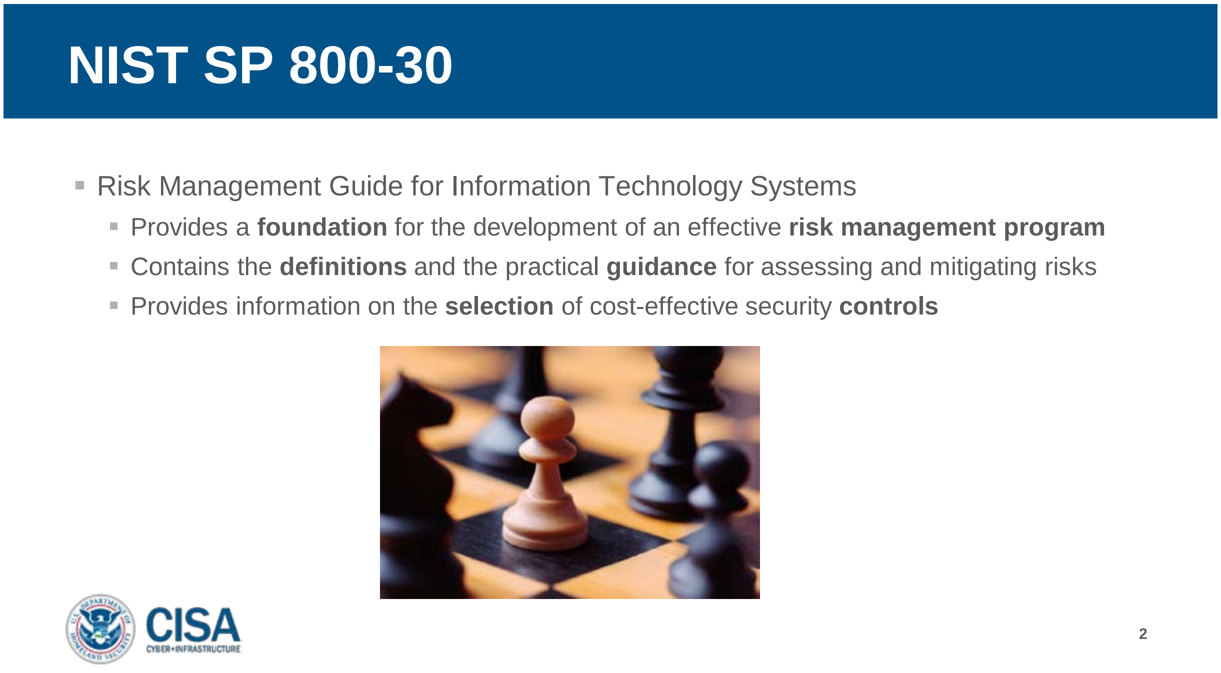

NIST SP 800-30

- Risk Management Guide for Information Technology Systems
  - Provides a **foundation** for the development of an effective **risk management program**
  - Contains the **definitions** and the practical **guidance** for assessing and mitigating risks
  - Provides information on the **selection** of cost-effective security **controls**

2

**002 Now, the National Institute of Standards and Department of Commerce, they've done a lot in terms of advancing the science of cybersecurity. I'm going to talk a little bit about some standards that are critical for cyber risk management. There's more than what I will cover in this presentation, but I'm giving you a high-level overview to you a basement of where you can start.

Let's start with NIST SP 800-30. This is a risk management guide for IT systems. It's a very good foundational document, and it'll actually give you a good sense for how to not only do some analysis of risks, but it'll also give you a good bit of idea of what controls you may have at your disposal that could be cost-effective for your organization in terms of addressing risk.

NIST SP 800-30: Risk Management

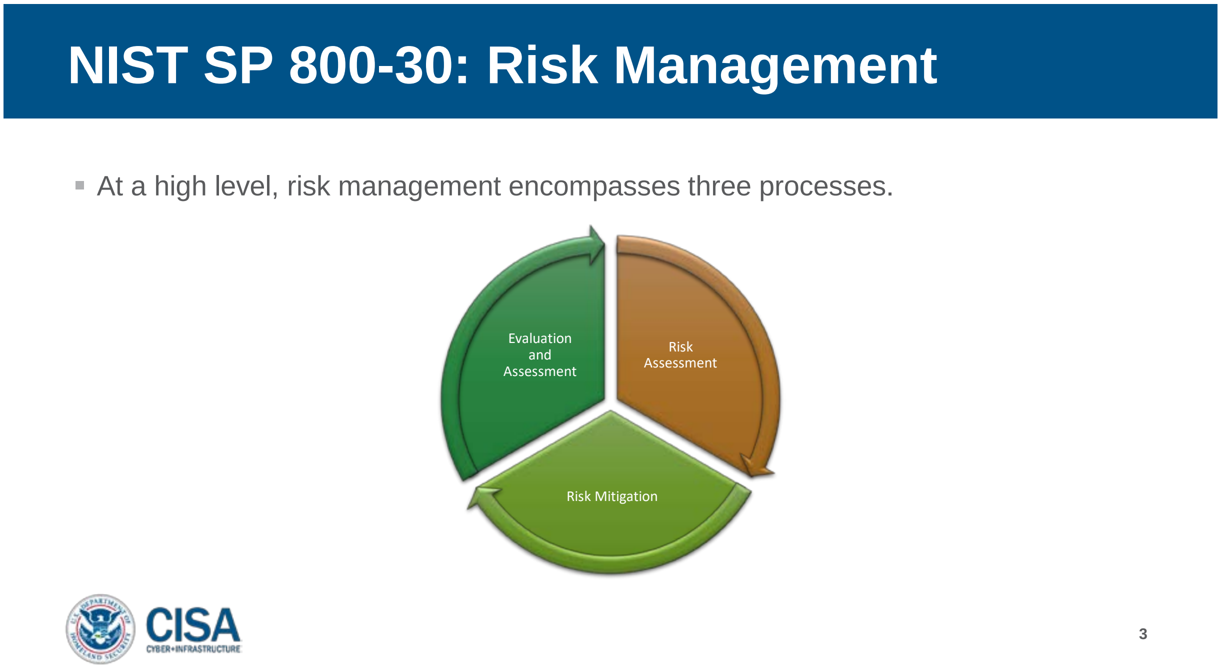

NIST SP 800-30: Risk Management

- At a high level, risk management encompasses three processes.

**003 And we can think about this at a high level. They break it out pretty good in terms of a three-step process, and each of these processes is layered with additional subset of actions to take within each of these overall elements. First we're going to do risk assessment, and then we're going to talk a little bit about risk mitigation, and then we're going to talk about how we go about evaluating and assessing and selecting risks toward the end.

Risk Assessment Steps Abstracted from SP 800-30

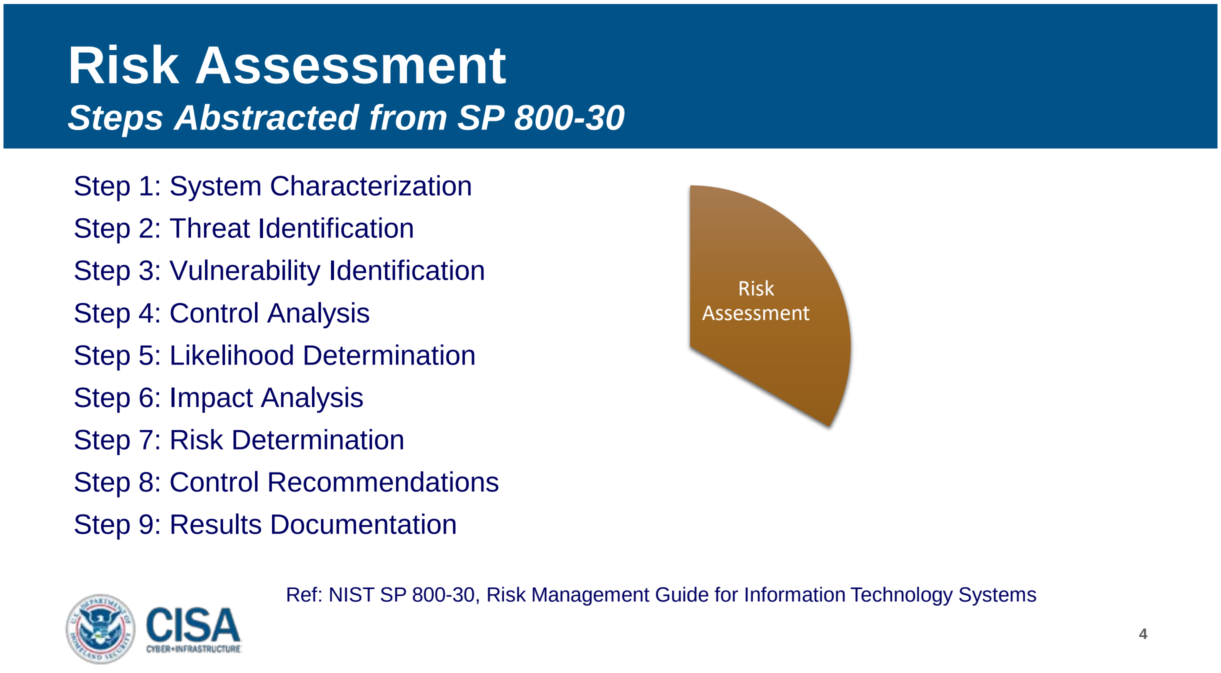

Risk Assessment

Steps Abstracted from SP 800-30

Step 1: System Characterization

Step 2: Threat Identification

Step 3: Vulnerability Identification

Step 4: Control Analysis

Step 5: Likelihood Determination

Step 6: Impact Analysis

Step 7: Risk Determination

Step 8: Control Recommendations

Step 9: Results Documentation

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

4

**004 Now, we're going to go through the first nine steps, and this is all dealing with risk assessment.

Step 1: System Characterization

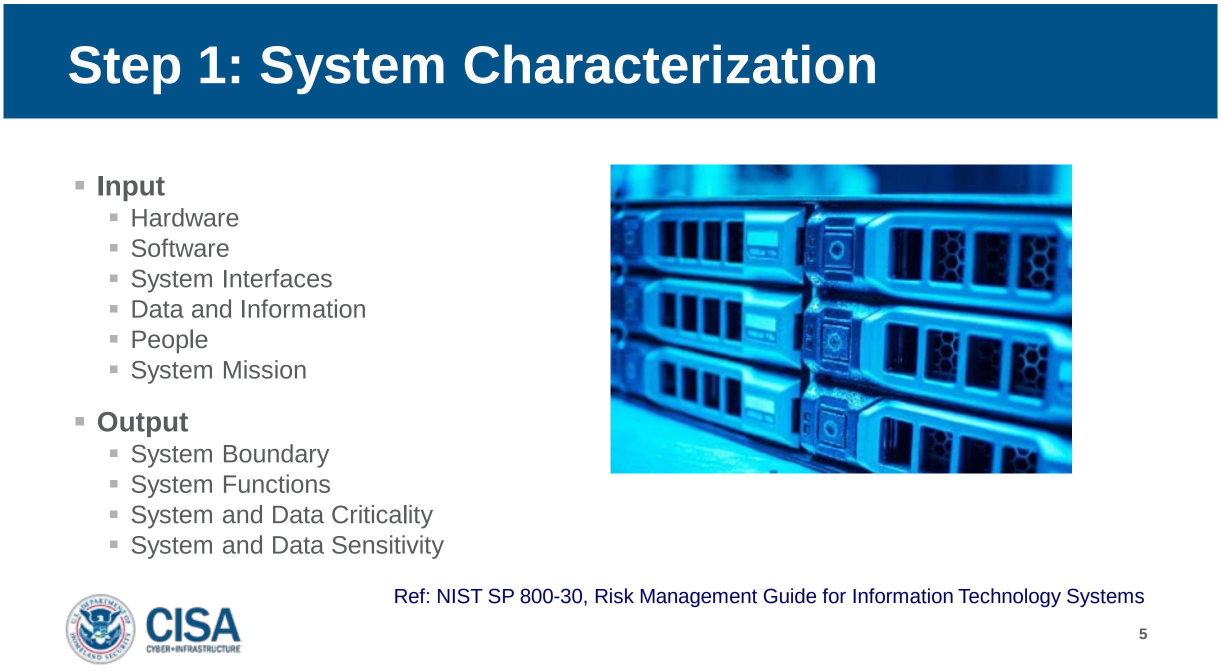

Step 1: System Characterization

- **Input**
  - Hardware
  - Software
  - System Interfaces
  - Data and Information
  - People
  - System Mission
- **Output**
  - System Boundary
  - System Functions
  - System and Data Criticality
  - System and Data Sensitivity

CISA
CYBER+INFRASTRUCTURE

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

5

**005 First where we want to start is with system characterization. Recall that in any given organization, your mission is to deliver some sort of critical service, or maybe you want to deliver some product of some sort, and it's going to take some kind of hardware or software that you're using, or maybe those systems are interfacing with others. Maybe you have data within those systems. Maybe, even more importantly, you have people who are operating them. And then you also have a mission that you're trying to accomplish. So we're going to use these elements, and hopefully what we get out of that analysis of those elements is we understand what our system boundaries are. Where is it that we have the absolute borders of what we're looking at trying to insulate from risks coming to fruition and actually impacting the enterprise negatively?

We also want to look at what are the functions. What's taking place within each of these elements that we have? We also want to look at the fact that we have data, obviously, but what we really want to look at more specifically is: Which data sets are more critical than others? Which are the crown jewels and which maybe are the kind that maybe you could do without if it were maybe by chance corrupted? Maybe if the confidentiality for some reason or another is no longer there, or maybe even if it's available to the organization.

And then we want to also talk about that confidentiality piece as to how sensitive that data could be. What's the impact if that data would be released to the greater public?

Step 2: Threat Identification

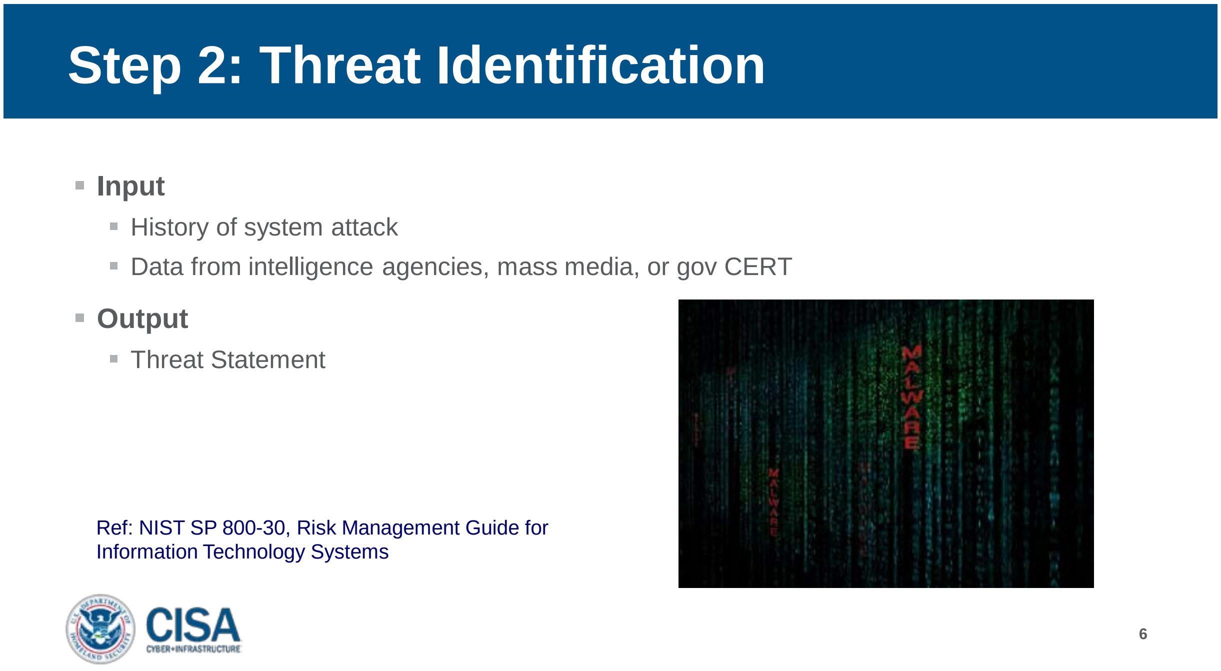

Step 2: Threat Identification

- **Input**
  - History of system attack
  - Data from intelligence agencies, mass media, or gov CERT
- **Output**
  - Threat Statement

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

MALWARE
MALWARE
CISA
CYBER+INFRASTRUCTURE
6

**006 So then we also want to think about threats. We want to think about the history of the system and we want to think about what are the actors wanting out of that system. Maybe it's that they want to interrupt your operations. Maybe it's that they want to understand or know elements of information that you have in that system. So you can get this threat identification from a bunch of different intelligence agencies; you can learn from mass media; there are other websites that you can google and find; you can even go to the CERT website and learn more about this.

Step 3: Vulnerability Identification

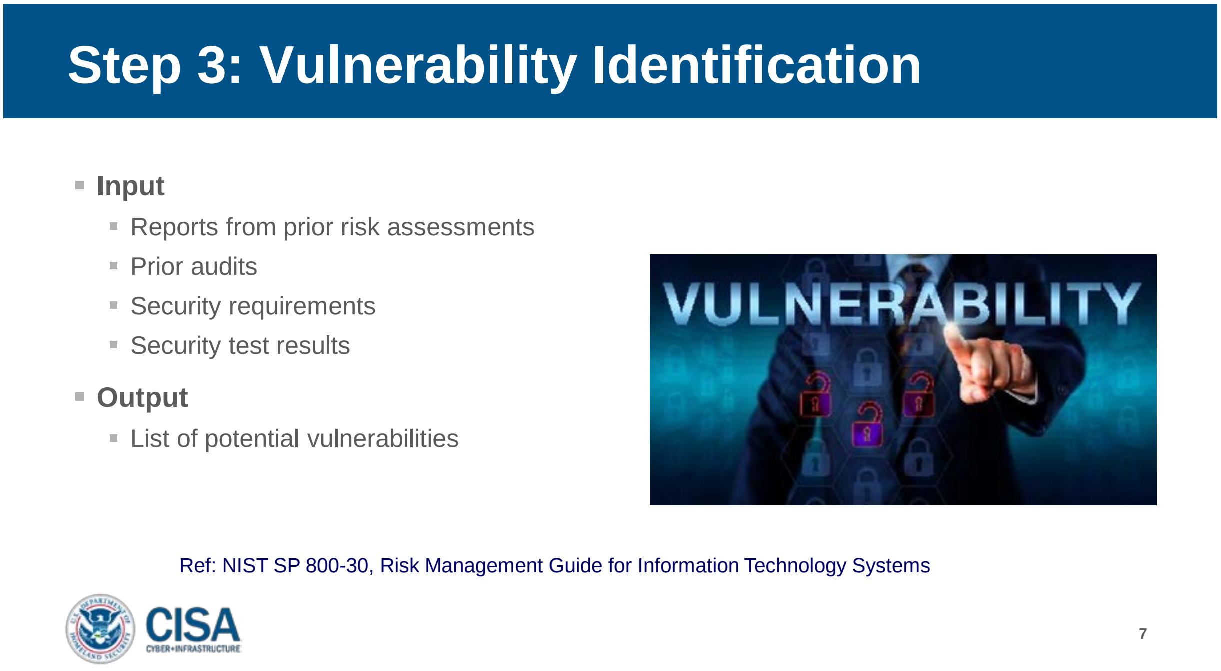

Step 3: Vulnerability Identification

- **Input**
  - Reports from prior risk assessments
  - Prior audits
  - Security requirements
  - Security test results
- **Output**
  - List of potential vulnerabilities

VULNERABILITY

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

CISA
CYBER+INFRASTRUCTURE

**007 And out of that, you're actually going to get a threat statement. Subsequent to that, you're going to think about vulnerabilities. You're going to want to think about what are the gaps in my systems that are going to let these threat actors in, and actually be successful.

So you want to think about reports from previous risk assessments you may have. More importantly, maybe you've had both internal and/or external audits that have taken place on your system that have given you maybe some sense or notion of what vulnerabilities may exist.

And maybe you have different security requirements, maybe at the organizational level, maybe even at a broader level, especially if you're a federal government entity, that is actually forcing you or driving you to have a certain set of security requirements. Maybe in the past too you've had tests done on your systems and you can use those test results. So those are all good inputs, and at the end of the day what you ultimately are looking for is a list of those potential vulnerabilities that may exist in your system.

Step 4: Control Analysis

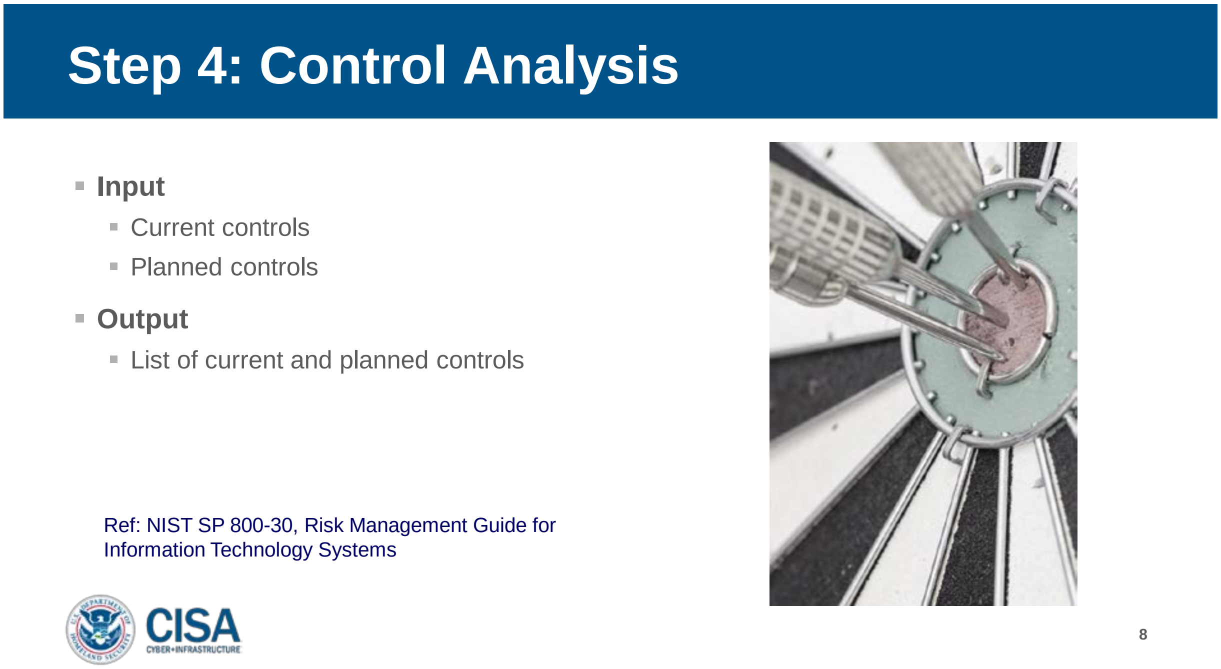

Step 4: Control Analysis

- **Input**
  - Current controls
  - Planned controls
- **Output**
  - List of current and planned controls

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

8

**008 Then what you're going to do is you're going to turn your mind to: Hey, now I understand how this risk may exist. I have a vulnerability, I have a threat actor, and maybe some idea even as to likelihood of them happening. How can I actually control it? What are the current controls that I have in place?

A good example here I like to think about, most organizations have fire extinguishers in their halls. You can go around and you can find fire alarms in your building, sprinkler systems, that kind of a thing. Those are current controls that are in place in case you have that dark day of the risk coming to fruition that you have a fire in your organization.

You may also have this notion of: What are my planned controls? What are the ones that I still need to put in place? Maybe I've procured them but I haven't implemented them. So there's all different states that these control sets could be in that you need to investigate and understand, identify, and document so that you can bring to the process.

Step 5: Likelihood Determination

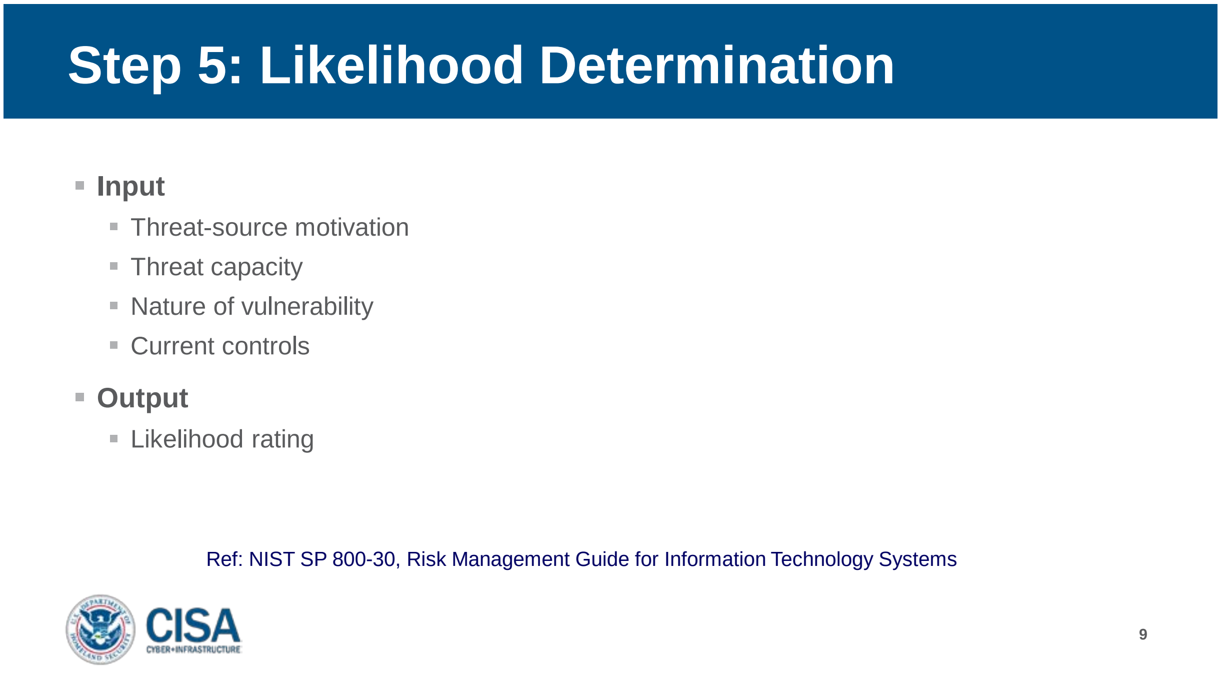

Step 5: Likelihood Determination

- **Input**
  - Threat-source motivation
  - Threat capacity
  - Nature of vulnerability
  - Current controls
- **Output**
  - Likelihood rating

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

9

**009 Now, one of the trickier parts of this is understanding the likelihood. So what we really want to do, the likelihood of this risk coming to fruition, is understand what are the motivations of these threat sources. What's their capability? What are they capable of? What's within their capacity?

A good example here is you think about a script kiddy, someone who's sitting in their basement and they just really want to hack your system so they can brag about it on the internet, or you may have a state actor who has a lot of resource, a lot of capacity, a lot of good, trained people that can do a lot of elaborate things in terms of techniques and trying to get into your system. That would actually come to play in terms of how likely they could be successful.

You also want to think about the nature of the vulnerabilities in your system, and you want to think about those controls, again, that you have in place. You kind of put that in a hopper and you all come out with this idea of a likelihood rating. Let me help you with this a little bit more, because we want to think about how to wrap your arms around likelihood.

Likelihood Rating Qualitative Ratings

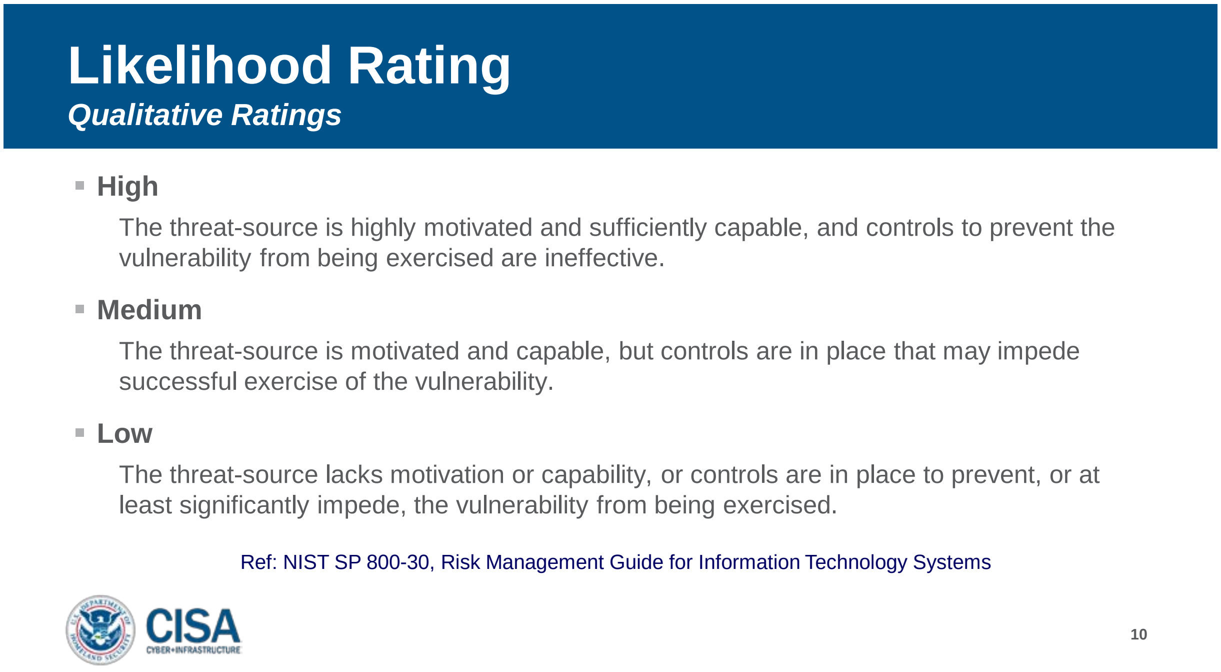

Likelihood Rating

Qualitative Ratings

- **High**

  The threat-source is highly motivated and sufficiently capable, and controls to prevent the vulnerability from being exercised are ineffective.

- **Medium**

  The threat-source is motivated and capable, but controls are in place that may impede successful exercise of the vulnerability.

- **Low**

  The threat-source lacks motivation or capability, or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised.

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

**010 So you're going to want to put ratings around this, and there's a way you could do this with qualitative ratings, and NIST 800-30 actually has some definitions around this. You could have high, medium or low characterization of how likely it is that these risks would take place. You could have a highly motivated actor that's really capable-- this could be that state actor, as I was talking about-- and you could know too that maybe you have limited or maybe even low-capability controls in place that could help you, and you could actually be rendered ineffective as an organization just as much as those controls could be ineffective for you to help keep your enterprise operating.

And from there, it just graduates down in terms of how the likelihood rating could come into place. For example, for medium, maybe the controls may be just impeded for a brief period of time. That way, the risk has a shorter window of actually coming into fruition, or maybe it's low, and now you've got it to a place where it's unlikely that the risk could ever occur.

Another Way to Think About Likelihood Ratings Quantitative and Functional Risk Appetite Statement

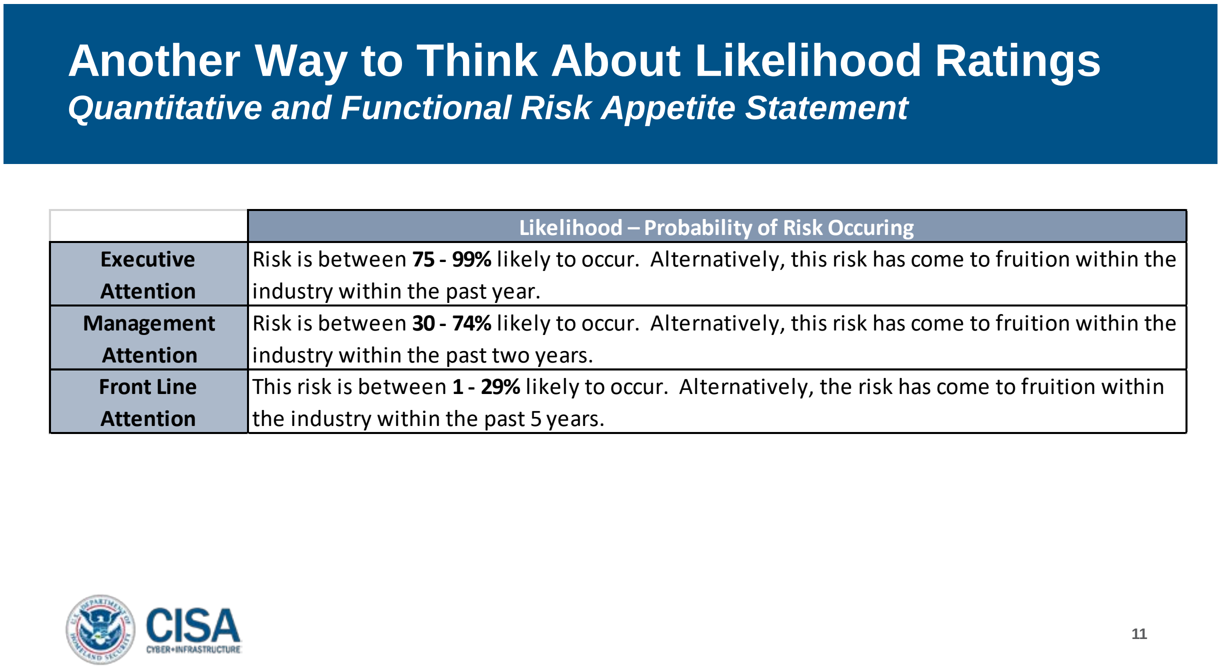

Another Way to Think About Likelihood Ratings

Quantitative and Functional Risk Appetite Statement

	Likelihood – Probability of Risk Occuring
Executive Attention	Risk is between **75 - 99%** likely to occur. Alternatively, this risk has come to fruition within the industry within the past year.
Management Attention	Risk is between **30 - 74%** likely to occur. Alternatively, this risk has come to fruition within the industry within the past two years.
Front Line Attention	This risk is between **1 - 29%** likely to occur. Alternatively, the risk has come to fruition within the industry within the past 5 years.

11

**011 So we could talk about this maybe in terms of appetite, just to make it a little bit clearer for you, and I like to think about risk appetite statements like I have in this table here. Some people are very quantitative in their mind. Engineers typically, they're good with numbers, right? Let's think of it that way. And you maybe dial it to a percentage of likelihood. I don't know, maybe it comes down to a coin toss, 50-50, right? Well, that would put me maybe in a band of maybe medium in terms of likelihood, and alongside with this appetite statement, what I'd like to point out is I've kind of put together a loose governance structure to associated with these likelihoods.

So if a risk is very likely, maybe that should shoot to the top in terms of having executive attention. Maybe if it's medium, maybe there's a possibility of it someday, but we know that it's certainly not today; it's only occurred in the industry maybe within the past couple years. Maybe that's at the management attention level. So there's ways you can think about this in terms of establishing appetite in an organization as well.

Step 6: Impact Analysis

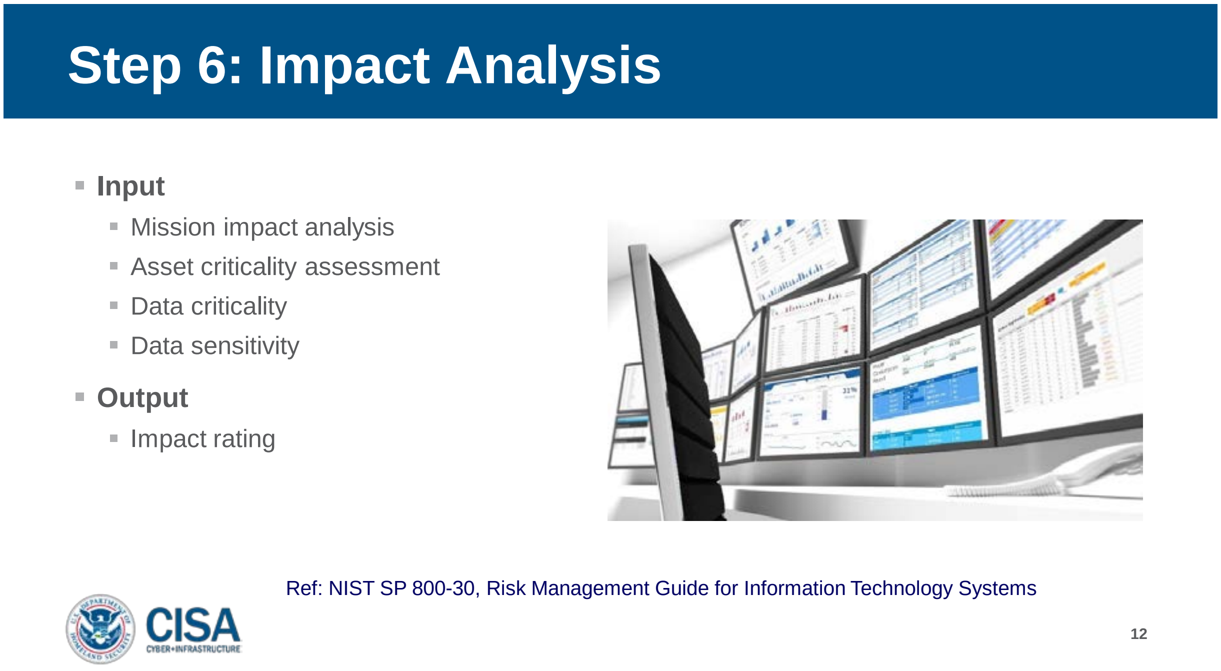

Step 6: Impact Analysis

- **Input**
  - Mission impact analysis
  - Asset criticality assessment
  - Data criticality
  - Data sensitivity
- **Output**
  - Impact rating

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

12

**012 Now there's other ways to think about risk too. Remember there's also the impact, and the inputs that we need in the process is we need to understand how the mission can be affected. We also have to understand too the assets, and with those assets, how critical they are to the organization so they can deliver those critical services. And once again, we've talked before about data criticality and sensitivity, and from that, we want to develop some sort of impact rating.

Impact Rating Qualitative Ratings

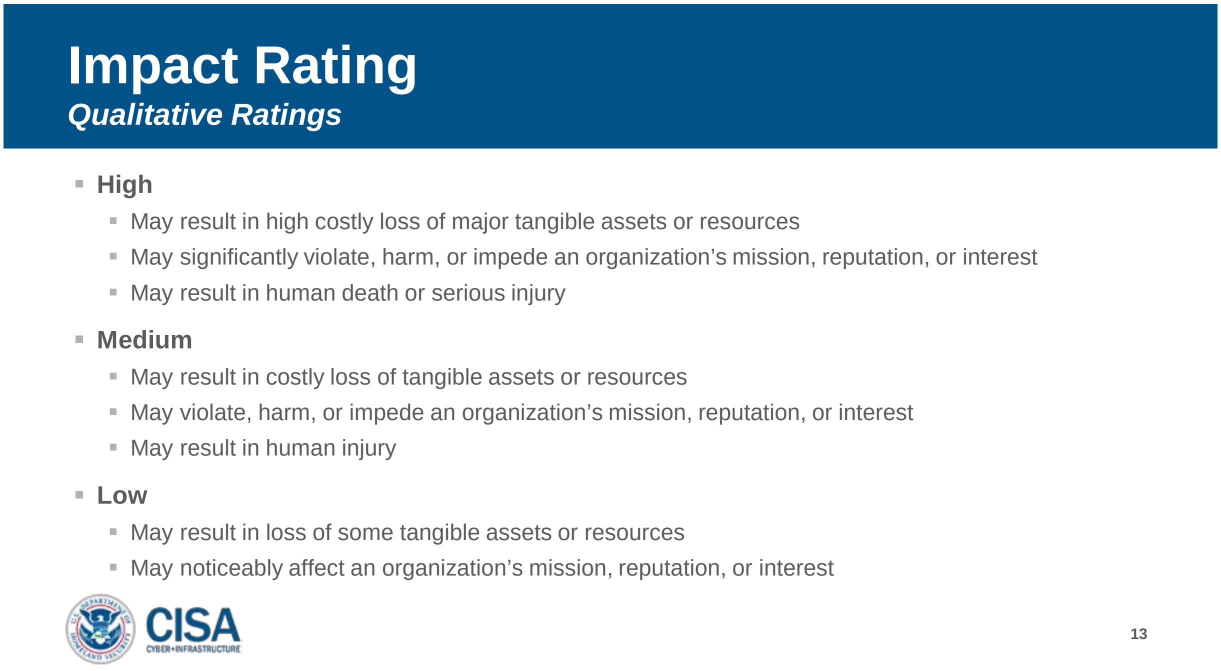

Impact Rating

Qualitative Ratings

- High
  - May result in high costly loss of major tangible assets or resources
  - May significantly violate, harm, or impede an organization’s mission, reputation, or interest
  - May result in human death or serious injury
- Medium
  - May result in costly loss of tangible assets or resources
  - May violate, harm, or impede an organization’s mission, reputation, or interest
  - May result in human injury
- Low
  - May result in loss of some tangible assets or resources
  - May noticeably affect an organization’s mission, reputation, or interest

CISA CYBER+INFRASTRUCTURE

13

**013 So let's go back to that model of the high, medium and low qualitative ratings that we could give. Now, this may vary from organization to organization-- it depends-- but this is kind of a high-level notion as to how that high, medium and low may play out.

So a high impact may result in maybe a very costly loss for your organization-- so much, in fact, that it could bring the organization to its knees, and maybe even make it so that the organization can no longer operate. Maybe it ruins the reputation of an enterprise to such an extent that you can no longer sell product anymore. Maybe you have a risk-- God forbid something happens- where you could have a death, or maybe a serious injury. That could be characterized as a high risk as well.

You can also have maybe a medium risk rating, where you could have maybe a costly loss; you could still operate, but it's going to be a long time before you can actually recover from that actual loss. Maybe people are harmed. Hopefully it's not to the fact that they would be maimed or it's a serious death.

So I think you're getting at least a sense here as to how these ratings could break out.

Let me go one step further and go back and go back to that appetite statement idea.

Another Way to Think About Impact Ratings Quantitative and Functional Risk Appetite Statement

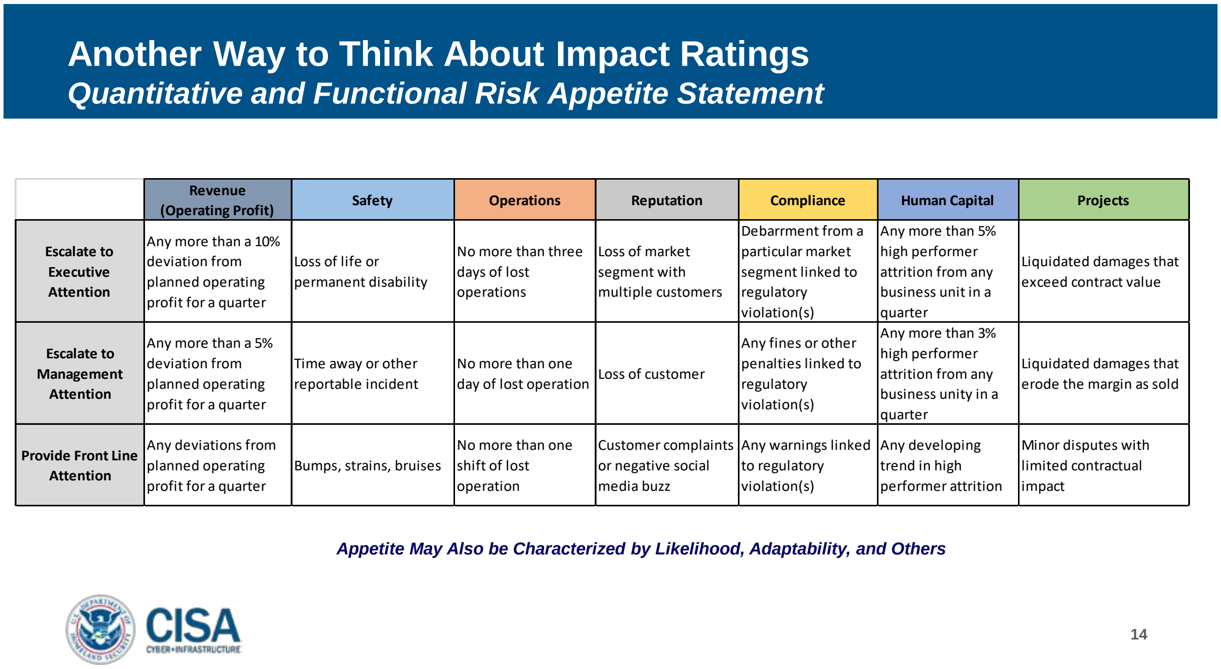

Another Way to Think About Impact Ratings

Quantitative and Functional Risk Appetite Statement

	Revenue (Operating Profit)	Safety	Operations	Reputation	Compliance	Human Capital	Projects
Escalate to Executive Attention	Any more than a 10% deviation from planned operating profit for a quarter	Loss of life or permanent disability	No more than three days of lost operations	Loss of market segment with multiple customers	Debarment from a particular market segment linked to regulatory violation(s)	Any more than 5% high performer attrition from any business unit in a quarter	Liquidated damages that exceed contract value
Escalate to Management Attention	Any more than a 5% deviation from planned operating profit for a quarter	Time away or other reportable incident	No more than one day of lost operation	Loss of customer	Any fines or other penalties linked to regulatory violation(s)	Any more than 3% high performer attrition from any business unity in a quarter	Liquidated damages that erode the margin as sold
Provide Front Line Attention	Any deviations from planned operating profit for a quarter	Bumps, strains, bruises	No more than one shift of lost operation	Customer complaints or negative social media buzz	Any warnings linked to regulatory violation(s)	Any developing trend in high performer attrition	Minor disputes with limited contractual impact

Appetite May Also be Characterized by Likelihood, Adaptability, and Others

CISA CYBER+INFRASTRUCTURE

14

**014 Now, with this particular appetite statement, it is not focused necessarily on likelihood; rather, it's focused on impacts. And you would see across the top, I have different categories that I have broken out here for how risk can impact an organization. Maybe it's a certain loss of revenue or resource. Maybe it's a safety issue. Maybe it's an operations issue, in terms of operation. Maybe you're really concerned in your enterprise about reputation, or complying to certain regulatory standards. Maybe it has to do with how you're managing your people, or how projects are taking place in your organization.

That top row that I'm talking about, the multicolored one there, can easily be born out of your organizational strategy. You look at your strategy statement and you understand what you're trying to achieve, and go ahead and fill in your own category it belongs. Safety, for example, is one that may be found universally through most organizations, and you can break it down in terms of who would care at what level in your organization to understand maybe what the level and degree of ratings should be. This is one example of how you could attack it. If you look at that safety category-- and notice that I have the loss of life as being the highest category-- and clearly there are going to be some people in the executive ranks that are going to be really concerned about that risk if it were to ever come to light.

Graduating down into that impact appetite for safety, you can see that maybe it's just time away, or maybe a reportable incident. Maybe that's at the management level of attention.

And it could be all the way down at the bottom of the chain, where maybe if it resulted in a bump, a strain, a bruise, maybe that's frontline supervisory-type concern. So yet another way to think about characterizing risk. I give that to you as an example.

Step 7: Risk Determination

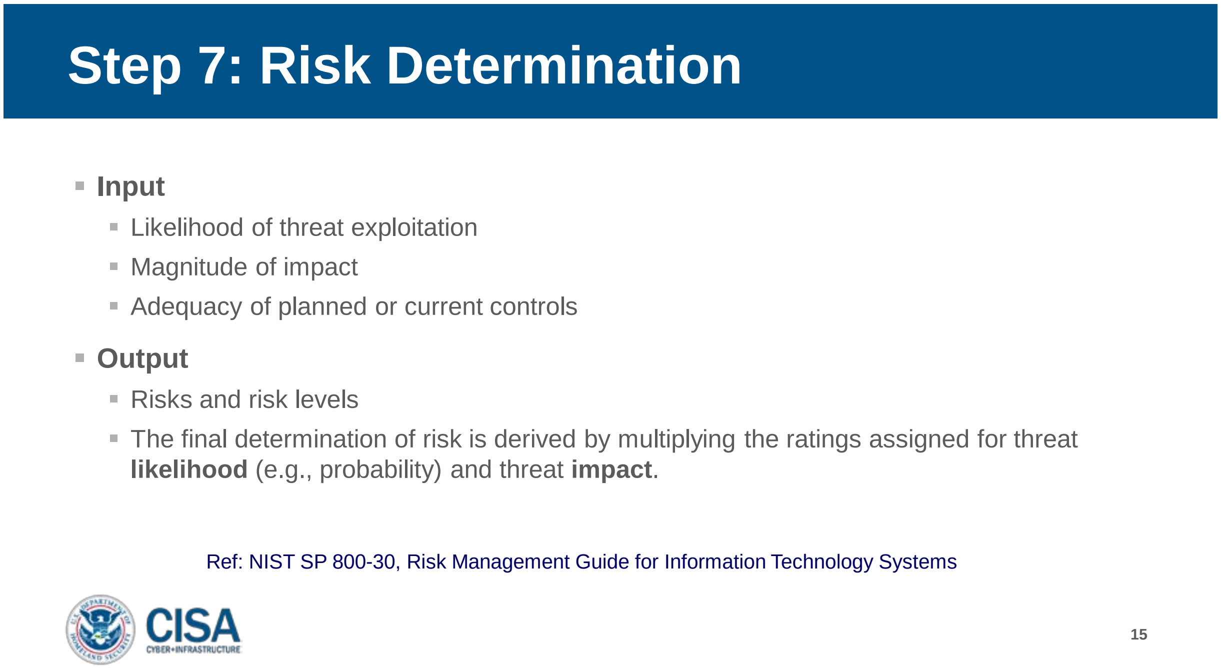

Step 7: Risk Determination

- **Input**
  - Likelihood of threat exploitation
  - Magnitude of impact
  - Adequacy of planned or current controls
- **Output**
  - Risks and risk levels
  - The final determination of risk is derived by multiplying the ratings assigned for threat **likelihood** (e.g., probability) and threat **impact**.

Ref: NIST SP 800-30, Risk Management Guide for Information Technology Systems

15

**015 So now we have all the elements together, and we really need to understand what the risk is, right? Maybe we understand the likelihood, we understand the threats, maybe even their capabilities. We understand maybe how it could impact the organization, and we also have documented and know what controls we have in place with respect to controlling these risks, and what we really want to get out of that is a sense for what the level of risk is in the organization, and what this is going to turn into is an understanding of what additional controls we may need, and also how do we prioritize these risks such that we know how to invest resources in a wise manner, because we don't have an infinite set of resources to answer every risk.

Step 8: Control Recommendations

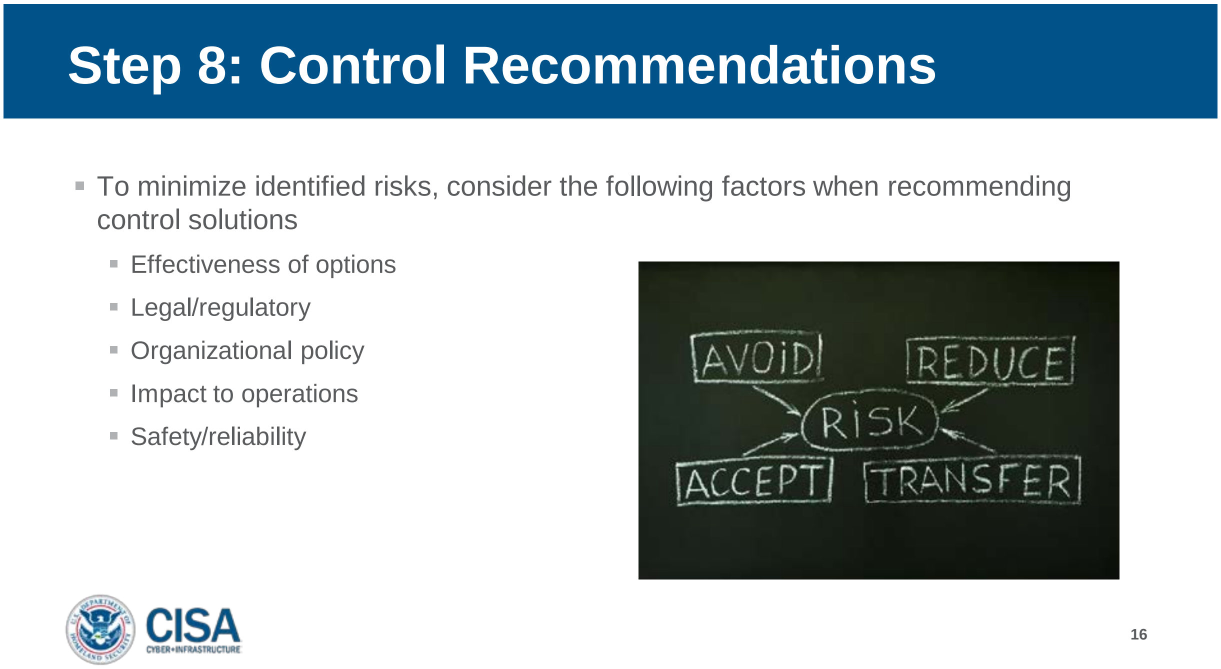

Step 8: Control Recommendations

- To minimize identified risks, consider the following factors when recommending control solutions
  - Effectiveness of options
  - Legal/regulatory
  - Organizational policy
  - Impact to operations
  - Safety/reliability

16

**016 So we want to think about, risk by risk, what are our control recommendations? What are the things that we can put into place to control each of these risks? Now, there are different strategies. Maybe what we decide to do is we actually maybe just accept a risk. That would mean that we just let it happen. Well, we don't necessarily just let it happen; there's some things that we want to think about if we're going to accept a risk. One is we want to document why we're just accepting that risk. Maybe the impact is to such a low level that really it doesn't matter if it were to happen.

Maybe it's so extreme we want to avoid the risk altogether, so we cease that operation or we don't even choose to embark upon a certain mission because the losses would just be too great, regardless of even maybe the likelihood of it happening.

Maybe we want to actually transfer that risk. Classically, in this case, in most organizations, a transference could be maybe you purchase insurance, so that way you're sharing that risk with another organization. Clearly you're paying the money to do that, but you've transferred away some of that problem.

Or maybe what you could do is take steps for mitigation. You can actually put controls in place that will maybe bring down the likelihood of it happening.

Now, there's a lot of drivers behind how you would select these strategies and these controls. Maybe there are regulatory drivers to it. Maybe you're required to have so many controls in place. For example, if I have a car and I'm selling cars to people who are going to be using them, they have to have seatbelts by law, because people have to wear them. This would be an example of a simple regulatory control for the risk that maybe that car would one day get in an accident, and it may save a life.

You may have organizational policies in place that will actually dictate what kind of controls that you need. Also you may want to step back and think about how my operations could be impacted, and that would be a big driver as to how you select those controls because, let's face it, once again, if you have a risk that comes to light and it actually brings your organization to its knees and it can't operate anymore, that may be a risk that comes to the top and demands a significant resource investment for control.

Safety-- yet another example.

Step 9: Results Documentation

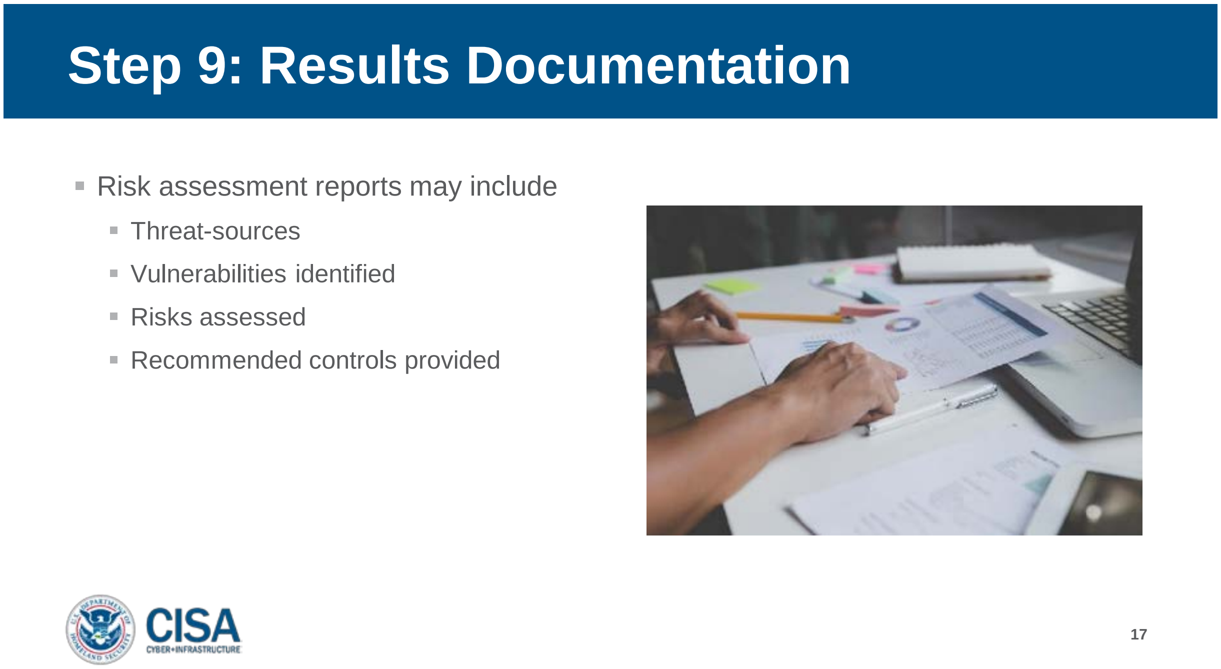

Step 9: Results Documentation

- Risk assessment reports may include
  - Threat-sources
  - Vulnerabilities identified
  - Risks assessed
  - Recommended controls provided

CISA
CYBER+INFRASTRUCTURE
17

**017 No matter what as you go through this process, you want to be sure to document what you've come up with. It goes without saying that we tend to, as people, float around organizations. You may in fact change jobs. Maybe you get promoted to a better job, and whoever comes in behind you is going to have to understand what that analysis was that took place. What assumptions were made? What were the threat sources at the time? It's not even necessarily the fact that you may leave that particular role. That risk may always be there no matter what controls you put in place, and iteratively you're to come back to visit it every year. You're going to really want to help yourself out and have that documentation in place so you can see what your mindset was for the assumptions that you made, for the vulnerabilities that you identified, for how you assess those risks. What was the process that you went through to actually come up with the results and the ideas and the decisions that you made?

Risk Mitigation – Steps 1 and 2

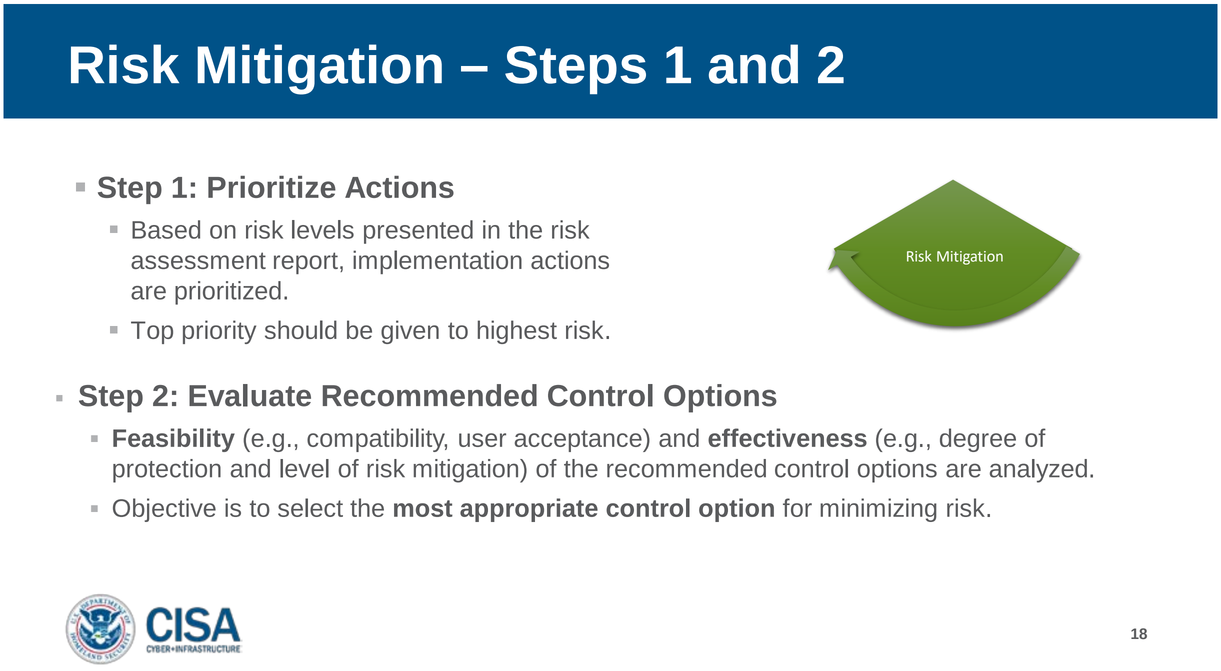

Risk Mitigation – Steps 1 and 2

- **Step 1: Prioritize Actions**
  - Based on risk levels presented in the risk assessment report, implementation actions are prioritized.
  - Top priority should be given to highest risk.
- **Step 2: Evaluate Recommended Control Options**
  - **Feasibility** (e.g., compatibility, user acceptance) and **effectiveness** (e.g., degree of protection and level of risk mitigation) of the recommended control options are analyzed.
  - Objective is to select the **most appropriate control option** for minimizing risk.

18

**018 So let's go down to the next larger step in that overall process, and it's going to talk a little bit more risk mitigation, and there's some steps here that we're going to go through to help you understand how you're going to go about implementing that response piece. Now remember, mitigation is a specific response. There are going to be actions related to it. Mitigation is nothing more than taking actions to control that risk.

Now, you're going to understand the risk levels of each risk by this point, because you've done your assessment, and you're going to have out of that a prioritization, and you want to understand which risks you're going to attack first. This may be a prioritization of actions with respect to one risk, or maybe it's a portfolio of risks and you have to understand the broader picture of all the actions needed to take place, and maybe there's a resource investment that needs to take place in terms of the time invested to take each of those actions. So there's a lot taking place in this first step.

You want to actually evaluate and look at the different control options you have too, because maybe it's not necessarily the actions that you take. How effective will those actions be at reducing that risk? Likewise, you may want to think about this in a different light. How effective are the people or resources you're leveraging against those risks in implementing those control options?

Risk Mitigation – Steps 3, 4, and 5

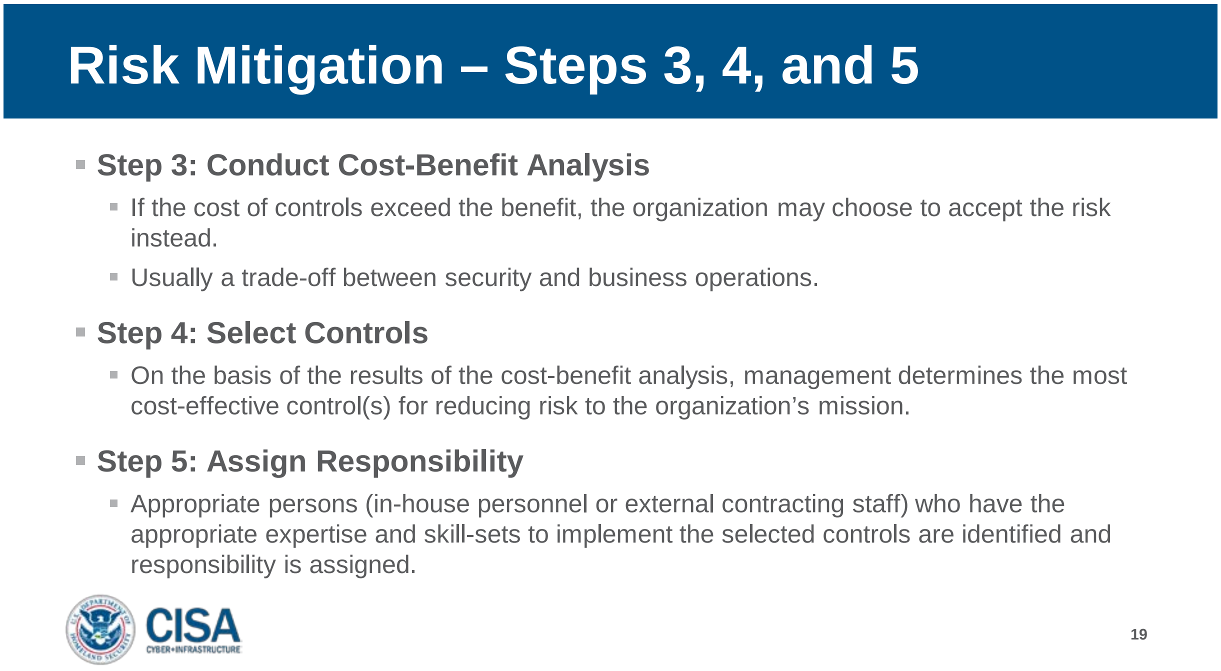

Risk Mitigation – Steps 3, 4, and 5

- **Step 3: Conduct Cost-Benefit Analysis**
  - If the cost of controls exceed the benefit, the organization may choose to accept the risk instead.
  - Usually a trade-off between security and business operations.
- **Step 4: Select Controls**
  - On the basis of the results of the cost-benefit analysis, management determines the most cost-effective control(s) for reducing risk to the organization’s mission.
- **Step 5: Assign Responsibility**
  - Appropriate persons (in-house personnel or external contracting staff) who have the appropriate expertise and skill-sets to implement the selected controls are identified and responsibility is assigned.

**019 So let's help you out a little bit here. In Step 3, maybe what you want to do to help get an understanding of prioritization and how these controls come about is doing a cost-benefit analysis. You may recall an example, if you've heard it before, is you don't want to build a million-dollar fence around a stack of pennies. This is the idea of doing a cost-benefit analysis for what you're investing in and how that control can actually help you in mitigating that risk.

You may come out with a list of certain controls, and maybe you can only afford several of them. So you're going to actually actively select what controls you want to put into place. That is a risk-based decision, and actually that is the heart of what we're trying to do here.

And then you're going to want to implement that control. To do that, you're going to actually assign responsibility to somebody, somebody who is a technician who understands how to put a certain control in place. Maybe I want to put a firewall in place in my system. Maybe I don't even have that certain person in house and I have to hire a vendor to do it. So you want to be mindful of this assignment of responsibility, not only in terms of who you select, but how are you going to monitor their performance? How are you going to make sure they're going to get done what they've promised to do? So it's going to want to be projectized, and you're going to have to follow maybe some metrics related to it. Maybe it's a "just do it" that happens right away and it's binary. It's, "Yes, it's done," "No, it's not done yet." Or maybe it's a process that could take months, or maybe even years to implement.

Risk Mitigation – Steps 6 and 7

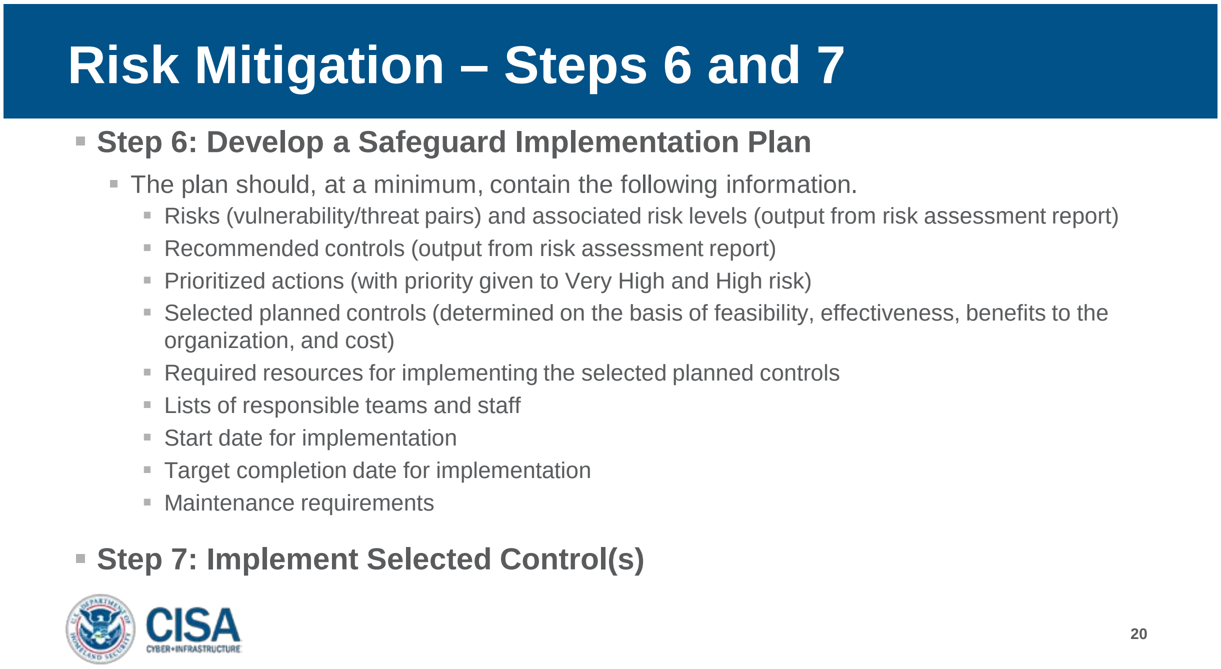

Risk Mitigation – Steps 6 and 7

- **Step 6: Develop a Safeguard Implementation Plan**
  - The plan should, at a minimum, contain the following information.
    - Risks (vulnerability/threat pairs) and associated risk levels (output from risk assessment report)
    - Recommended controls (output from risk assessment report)
    - Prioritized actions (with priority given to Very High and High risk)
    - Selected planned controls (determined on the basis of feasibility, effectiveness, benefits to the organization, and cost)
    - Required resources for implementing the selected planned controls
    - Lists of responsible teams and staff
    - Start date for implementation
    - Target completion date for implementation
    - Maintenance requirements
- **Step 7: Implement Selected Control(s)**

CISA

20

**020 This all comes down to how you have an implementation plan in place. So you're going to have to understand what the risk is, you're going to have those recommended controls, and you're going to do a lot to monitor these teams and make sure that they're actually doing what they need to put into place.

You also want to think about, once you have these controls in place, how to phase it into an organization. There may be some change management involved here. So you're going to have some target completion dates and that change management plan is going to have to lay out how that organization is going to transition to this new control, and once the control is in place, you may even have to think about how you're going to maintain it once it is in place. Maybe there's certain configurations that change over time. Maybe there are updates that need to take place. Who's going to do it? How? And with what resources?

Evaluation and Assessment

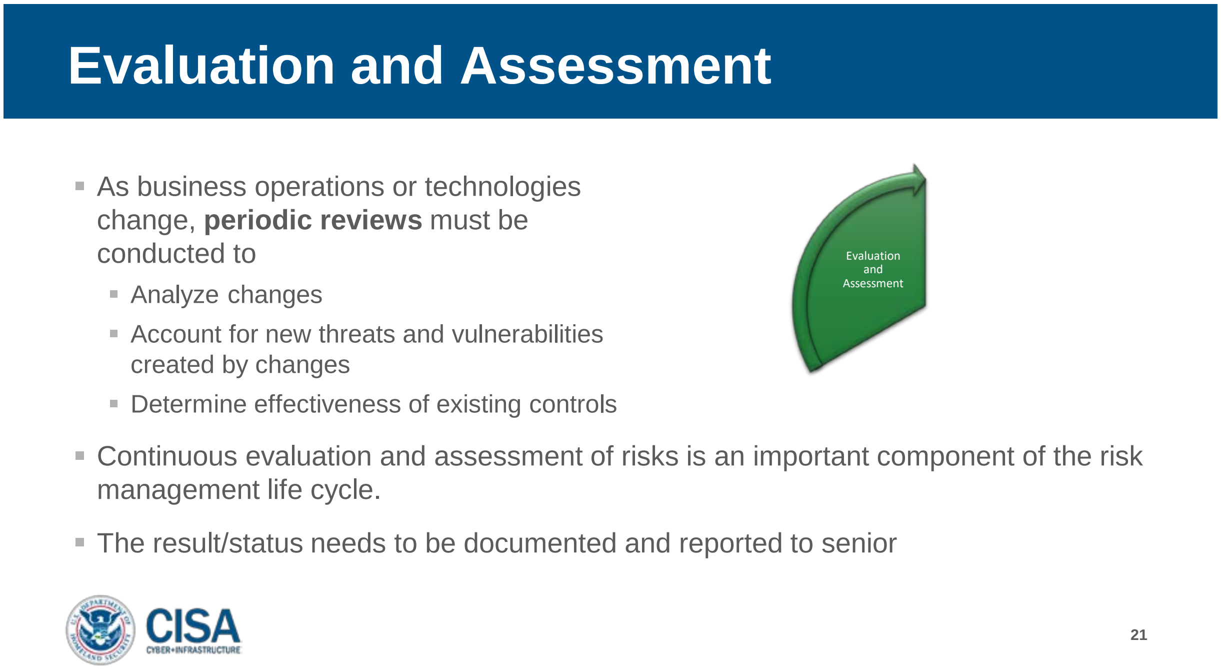

Evaluation and Assessment

- As business operations or technologies change, **periodic reviews** must be conducted to
  - Analyze changes
  - Account for new threats and vulnerabilities created by changes
  - Determine effectiveness of existing controls
- Continuous evaluation and assessment of risks is an important component of the risk management life cycle.
- The result/status needs to be documented and reported to senior

21

**021 Then, once those controls are implemented, you have met the bounds of that certain mitigation step.

So now let's turn to evaluating and assessing how those controls are doing for you. You're going to do this by analyzing changes. You're going to look for new threats and vulnerabilities that may be on the horizon and how that's impacting the control sets that you have in place, and you also want to look at how effective those existing controls are in their performance. Are they doing what they said they would do? Are they delivering as to how you thought those controls would perform? And this is a continuous process. It doesn't just end with one test. You're going to want to iteratively keep looking back and evaluating and checking back to make sure that that risk is being controlled like you thought it should.

NIST SP 800-39

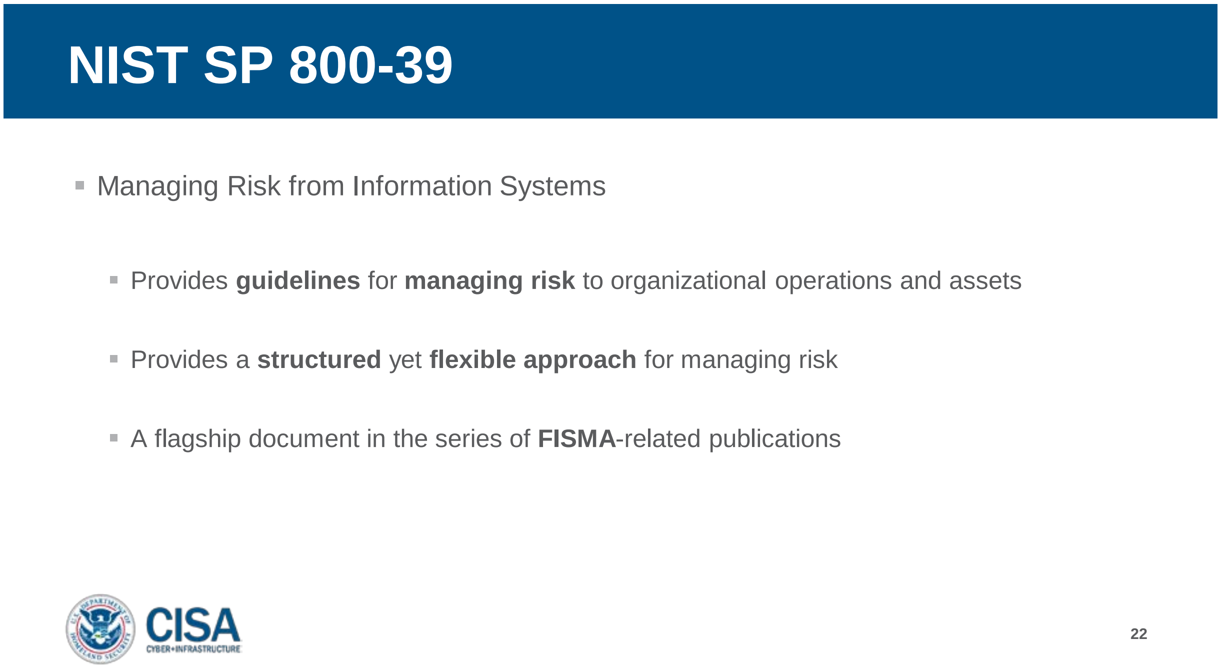

NIST SP 800-39

- Managing Risk from Information Systems
  - Provides **guidelines** for **managing risk** to organizational operations and assets
  - Provides a **structured** yet **flexible approach** for managing risk
  - A flagship document in the series of **FISMA**-related publications

CISA
CYBER+INFRASTRUCTURE

22

**022 In NIST 800-39, they're really good about this, and there's the Federal Information System Management Act that has a number of publications in this regard that dictate how risk will be managed in information systems. I highly recommend going to look at this 800-39. It has a lot of good guidelines to put you through the process of managing this risk, and it gives you a structured and flexible approach to doing so.

NIST SP 800-39: Tiers of Risk Management

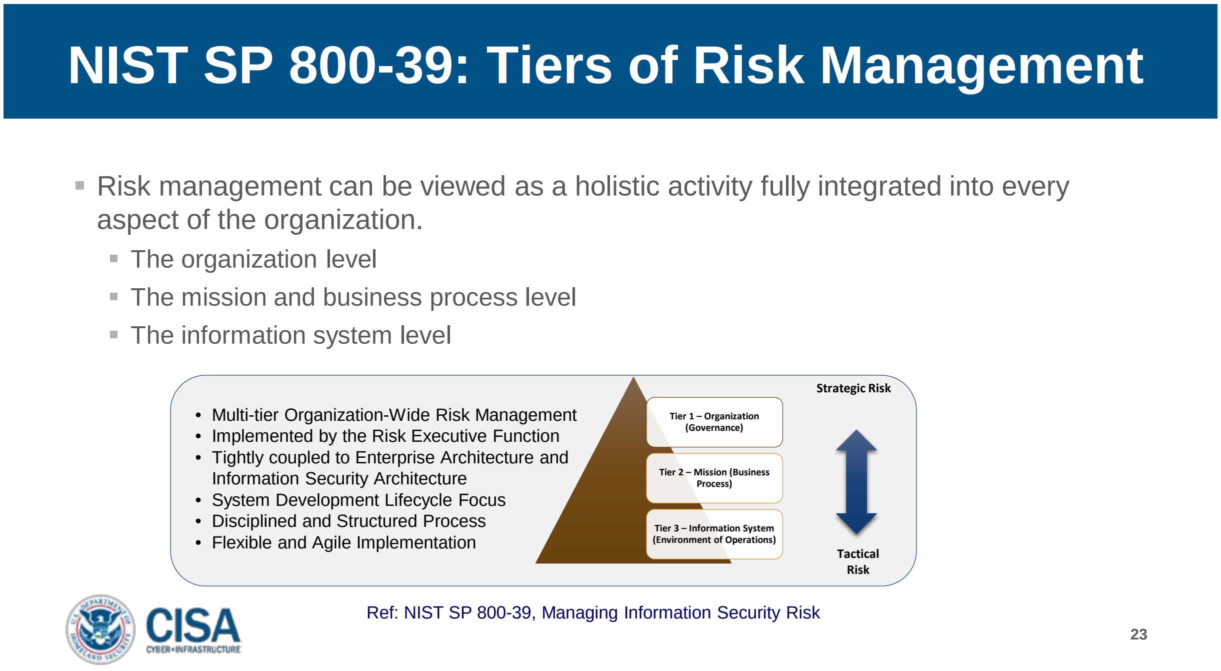

NIST SP 800-39: Tiers of Risk Management

- Risk management can be viewed as a holistic activity fully integrated into every aspect of the organization.
  - The organization level
  - The mission and business process level
  - The information system level

- Multi-tier Organization-Wide Risk Management
- Implemented by the Risk Executive Function
- Tightly coupled to Enterprise Architecture and Information Security Architecture
- System Development Lifecycle Focus
- Disciplined and Structured Process
- Flexible and Agile Implementation

Tier 1 – Organization (Governance)

Tier 2 – Mission (Business Process)

Tier 3 – Information System (Environment of Operations)

Strategic Risk

Tactical Risk

Ref: NIST SP 800-39, Managing Information Security Risk

23

**023 It also does a good job at breaking out how risks should be managed across an enterprise. You may think of your enterprise broken up in tiers. Truly, at the very top of the structure, you're going to have strategic risks that are going to need to be considered. These are the ones that are long-range and may have a large global impact to your organization.

You may also have missions that are taking place underneath that organizational layer, at the high level, and you have processes that support those missions getting accomplished. You also have assets related to it. You may want to think about that at yet another tier, where risks exist as well.

And then at the lowest level, you may have tactical risk-- so for things that are taking place on any given day. For example, maybe I'm operating a manufacturing operation. There is a tactical risk that maybe someone-- hopefully does not-- but they could get injured while they're doing work. That would be a risk at the tactical level. Same idea could be said about information systems that maybe they get hacked in a day. Maybe there's a phishing email that gets opened. That would be a tactical risk as well.

NIST SP 800-39: Process Applied

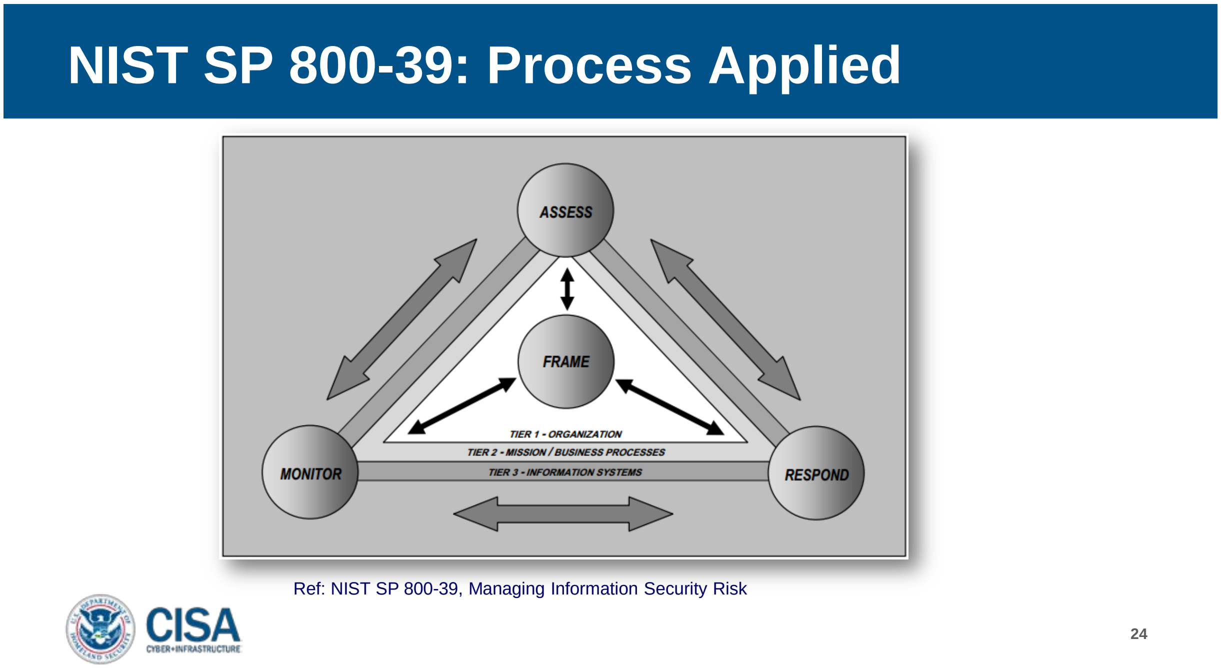

NIST SP 800-39: Process Applied

Ref: NIST SP 800-39, Managing Information Security Risk

24

**024 So 800-39 breaks this out into a process where we're continually assessing, we're responding, and we're monitoring those risks, and it all takes place within the frame of how those risks actually exist, and it sees it through the lens of these tiers. This is a graphic from 800-39 that kind of spells that out for you and makes it a little bit more solid in your mind, lets it jell a little bit better.

NIST SP 800-39: Risk Framing

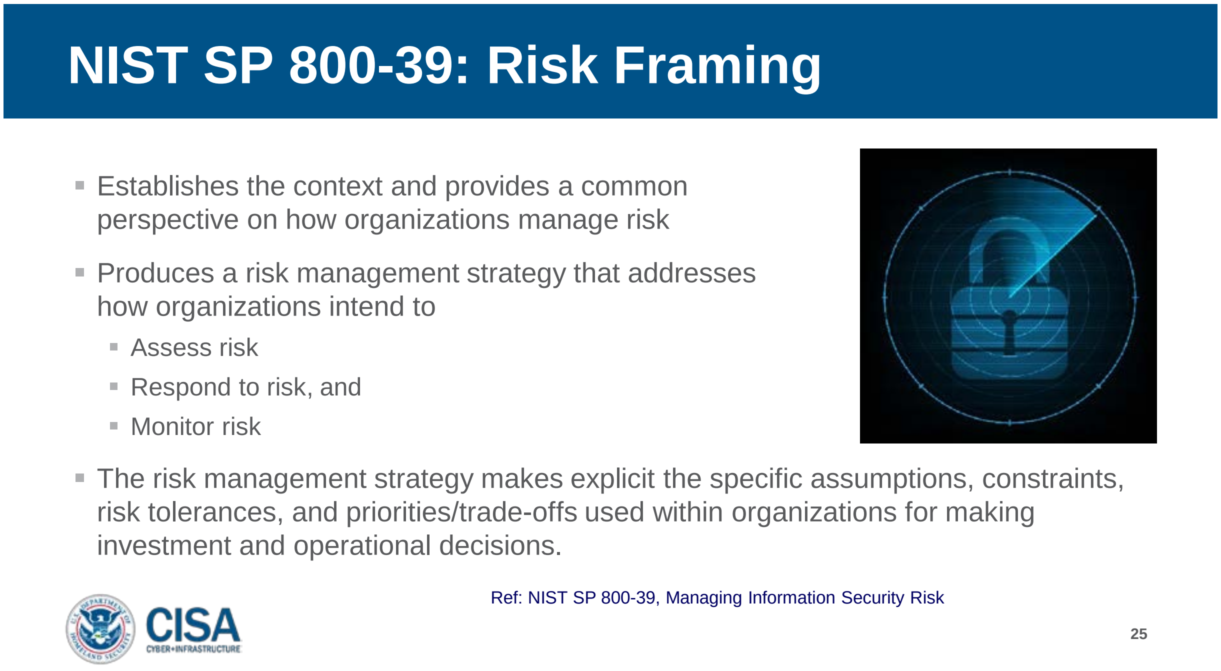

NIST SP 800-39: Risk Framing

- Establishes the context and provides a common perspective on how organizations manage risk
- Produces a risk management strategy that addresses how organizations intend to
  - Assess risk
  - Respond to risk, and
  - Monitor risk
- The risk management strategy makes explicit the specific assumptions, constraints, risk tolerances, and priorities/trade-offs used within organizations for making investment and operational decisions.

Ref: NIST SP 800-39, Managing Information Security Risk

**025 Let's talk a little bit about that risk framing. It's the idea of establishing a context around a risk and how it may actually come to light. You're going to want to look at how they're going to actually assess risk, how you're going to respond to it, and how you're going to monitor it; and to have that understanding of how you're going to do those steps, you're going to want to know the greater organization that the risks exists in, you're going to want to know the mission it exists in, and that's all to context-- that's speaking to the context.

NIST SP 800-39: Risk Monitoring

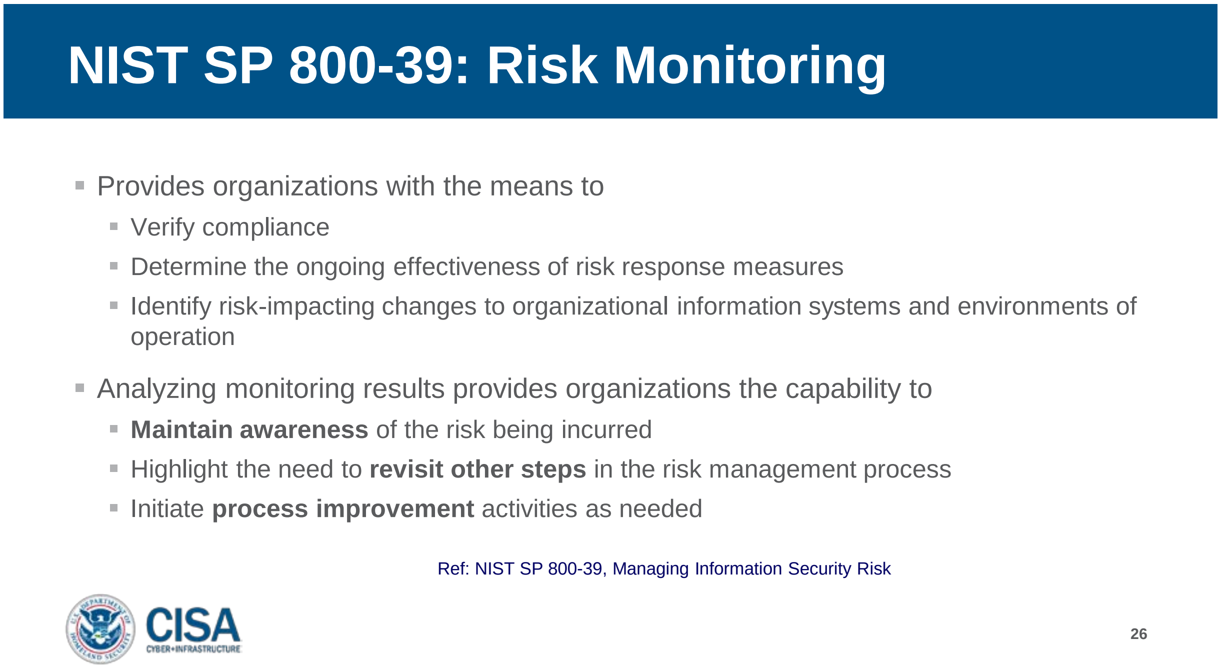

NIST SP 800-39: Risk Monitoring

- Provides organizations with the means to
  - Verify compliance
  - Determine the ongoing effectiveness of risk response measures
  - Identify risk-impacting changes to organizational information systems and environments of operation
- Analyzing monitoring results provides organizations the capability to
  - **Maintain awareness** of the risk being incurred
  - Highlight the need to **revisit other steps** in the risk management process
  - Initiate **process improvement** activities as needed

Ref: NIST SP 800-39, Managing Information Security Risk

26

**026 Once you understand that, now you're going to know how to monitor your risk. You're going to need to know if you're doing it to be compliant with maybe a certain regulation. There may be this notion that there's effectiveness of the response measures that you put in place. Maybe you have certain metrics that you've put in place to measure that effectiveness. It's important to understand what those metrics are. And actually, you want to go another step. Think about who's monitoring those metrics. Who's compiling the data and information? What are they doing with it, and who are they reporting to?

You may also want to think about how you're going to monitor those results and how you're going to deliver them such that you can maintain the capability. How do I know if the risk has actually occurred? You may have to go about establishing key risk indicators. These are things that may happen that, if they do, it would be indicative of the risk actually coming to fruition. You want to be aware of that, and you want others in your organization to be aware of it too.

And there's also this idea too that once you have these processes in place, you may need to iteratively improve upon them to make them better.

NIST SP 800-39: Risk Response

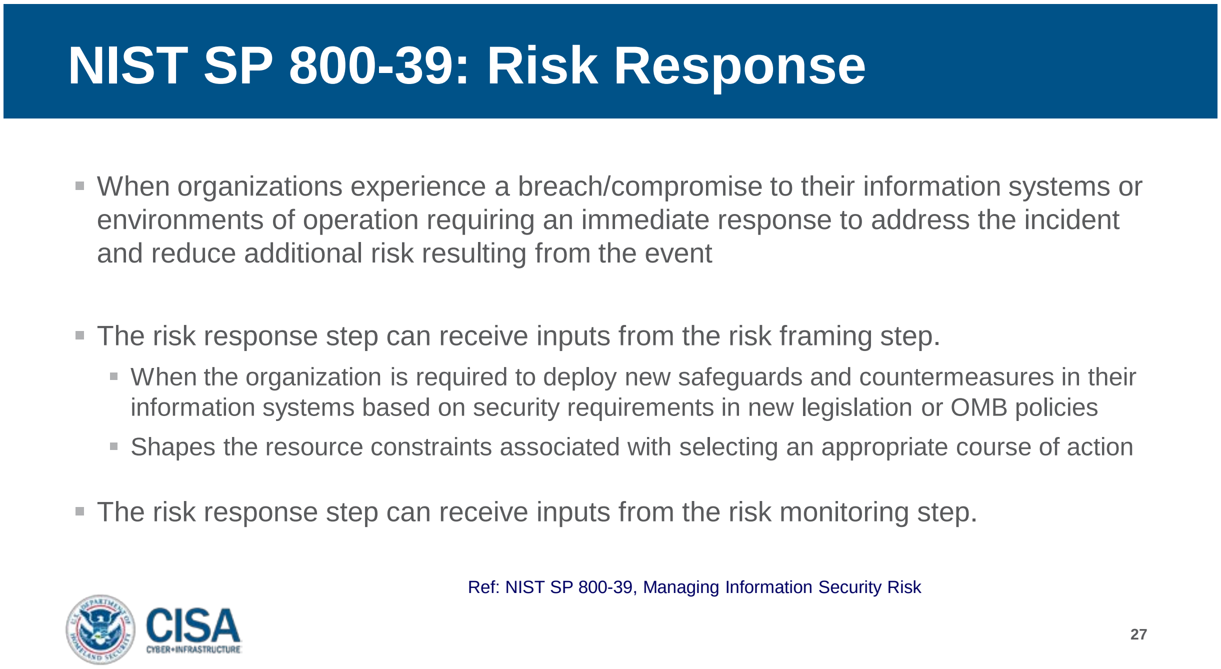

NIST SP 800-39: Risk Response

- When organizations experience a breach/compromise to their information systems or environments of operation requiring an immediate response to address the incident and reduce additional risk resulting from the event
- The risk response step can receive inputs from the risk framing step.
  - When the organization is required to deploy new safeguards and countermeasures in their information systems based on security requirements in new legislation or OMB policies
  - Shapes the resource constraints associated with selecting an appropriate course of action
- The risk response step can receive inputs from the risk monitoring step.

Ref: NIST SP 800-39, Managing Information Security Risk

**027 And you're actually to the response part now too. So this is that time when that cold, dark day comes, when the risk has come to fruition, and you have to actually respond to the risk actually taking place. This is where incident management will come into place, and it's good to have a step-by-step thought as to how this response will go down. What are the things that are specifically going to need to be done so that way you can limit the damage that is taking place? So you may have to put new safeguards in place or countermeasures to make sure that you're limiting that damage that takes place, and some of it may even be required in a regulatory sense.

It also-- this particular step-- would actually start to speak to how your resources are actually going to be applied, and you're going to find very quickly that some of those resources could be constrained.

A good example is if you have a risk that shuts down your entire system-- let's say ransomware-- and let's say it happens late at night, and your security and operations center maybe only has one or two people that are there in the evening. Well, to be honest with you, they may have a phone tree where they're going to call a whole lot of people to come support them. So they're constrained at the moment to have any more better of a response than actually just getting help. You're going to think that through.

NIST SP 800-37

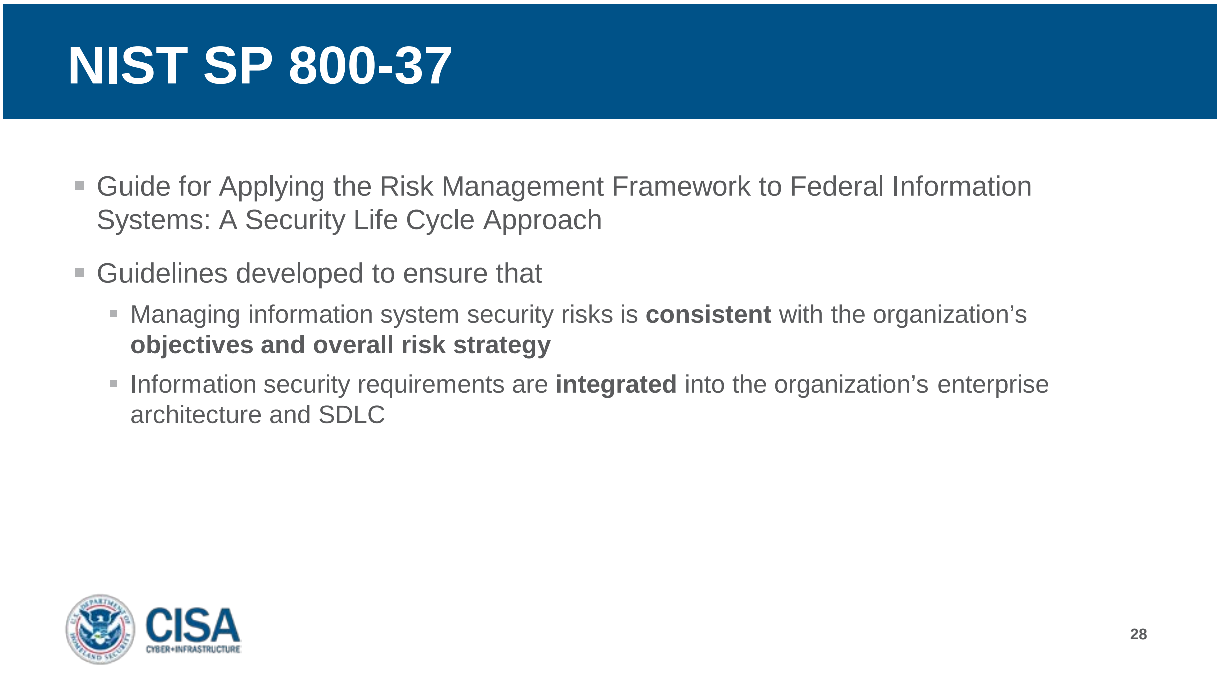

NIST SP 800-37

- Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach
- Guidelines developed to ensure that
  - Managing information system security risks is **consistent** with the organization’s **objectives and overall risk strategy**
  - Information security requirements are **integrated** into the organization’s enterprise architecture and SDLC

28

**028 So there's another standard out there, and it's NIST 800-37, and this is the guide for applying a risk management framework, also known as the RMF, for federal information systems. Yet another good document. It actually goes through the software development lifecycle, and it treats it in an integrated sense so that way you understand how you may or may not be introducing risk in your enterprise as you're bringing new assets into the organization, and it aligns it strongly with what the objectives of the organization are.

Risk Management Framework

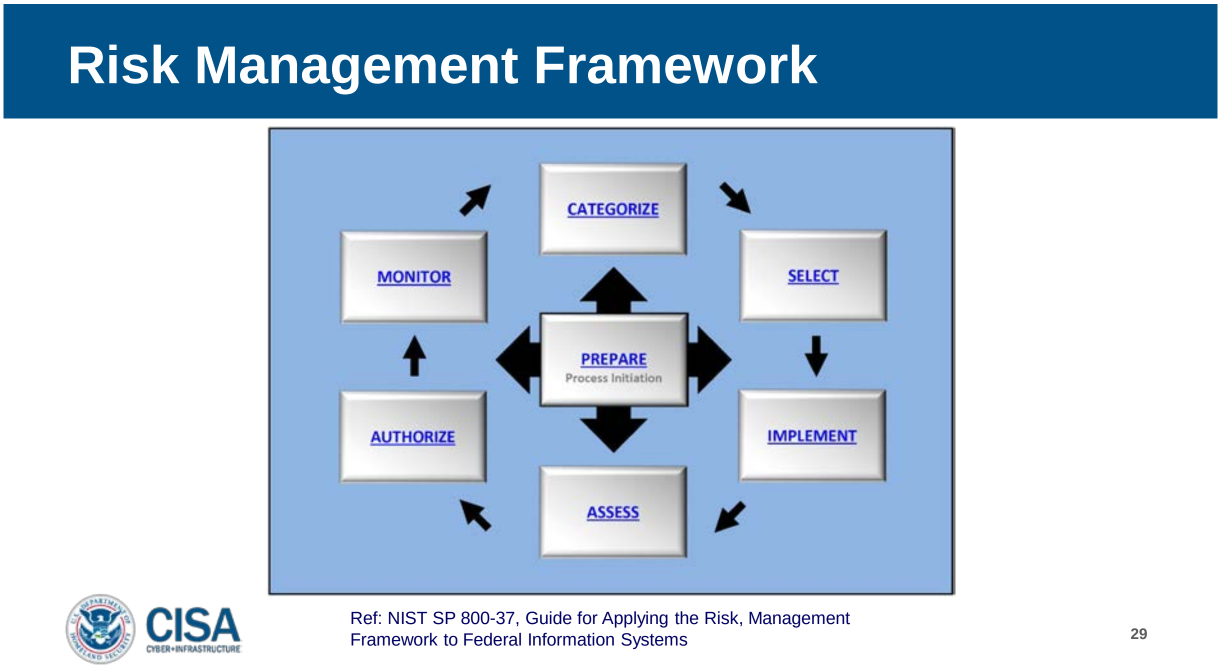

Risk Management Framework

PREPARE
Process Initiation

CATEGORIZE

SELECT

IMPLEMENT

ASSESS

AUTHORIZE

MONITOR

Ref: NIST SP 800-37, Guide for Applying the Risk, Management Framework to Federal Information Systems

29

**029 It was recently updated, and Here's the basic steps, And I'm going to go through each of these at a high level, but I highly encourage you to go review the document, because first you want to be able to categorize, and not just categorize risks yet. What you really find out in that step is you're actually categorizing your assets and your services. With that, then you identify what your risks are, so that way you understand what controls you're going to select and how you're going to go about implementing them and how you're going to assess their effectiveness.

And then there's this notion of, "Okay, now that I understand what the asset is, I understand what the risks are related to it, I have controls in place. I'm going to actually authorize that that system or that asset may come into operation in the enterprise, and then I'm going to monitor it." Now, all this is done within the context of how you prepare to go about this process, which is a new step in the RMF. I highly encourage you go look at it and review it and find out what's taking place in that step.

Notices

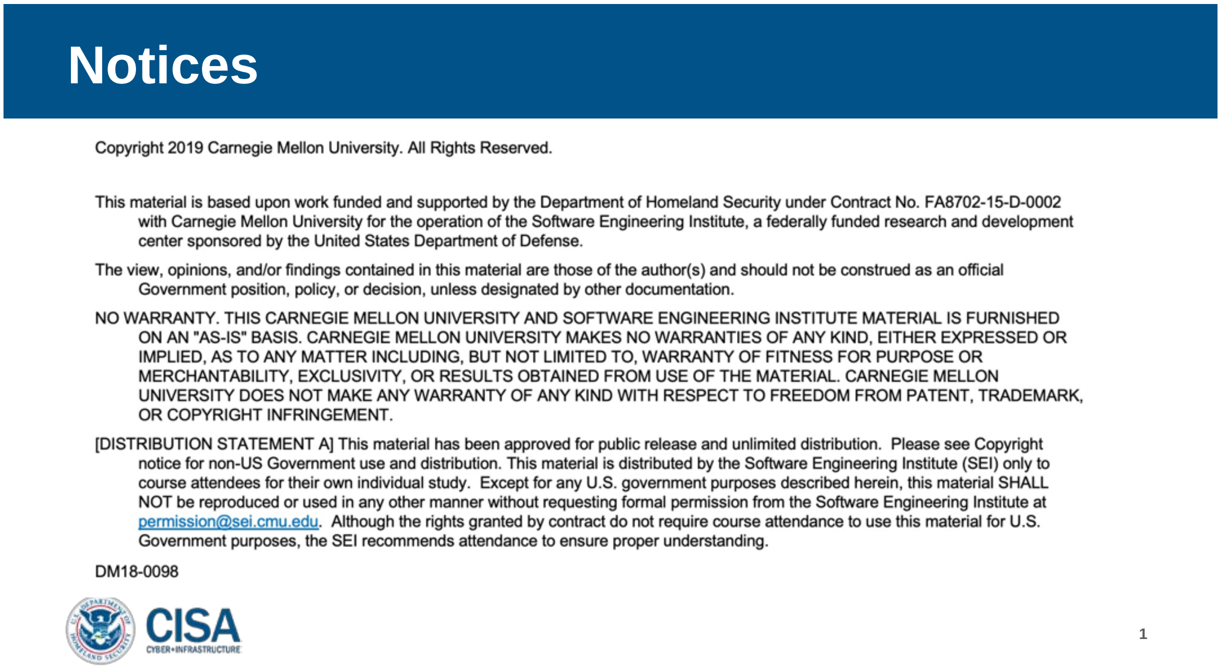

Notices

This material is based upon work funded and supported by the Department of Homeland Security under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

The view, opinions, and/or findings contained in this material are those of the author(s) and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

DM18-0098