Practical Cyber Security for Cyber Security Practitioners

Prof. Sandeep Kumar Shukla

Department of Computer Science and Engineering

Indian Institute of Technology, Kanpur

Lecture 20

So, good news and bad news. Good news is that we have new project teaming announced on Canvas. So, now we have 14 teams. The other good news is that there would not be any term project; there will be 3 more homeworks. So, we were initially supposed to have four homeworks. We had two homeworks done already, and two more were due. Instead of having a team project, I looked at last year's team projects, and I did not like what I saw. So, I said this is not going to be a very useful project. So, I have decided to make it three projects. The other good news is now you can do these 3 assignments or projects, whatever, in groups of 5 or 6, rather than groups of 3.

So, we will have 14 groups, as you can find on Canvas. The bad news is what that means: your final grade will be highly dependent on your exams, right? And as you know, the exam results were pretty abysmal, as expected, because you can see that out of 83 students, here I have maybe 13 students. The questions are such that if you do not listen to the lectures, you would not be able to answer anything, right? Just by reading the slides. Having said that, let us move on, and also let me show you where we are with respect to the curriculum. So, we are now—so I missed a week of classes, and then Nanda did 2 weeks of classes. So, I might have to do extra classes later in the semester, maybe after the midterm break, to cover some of these.

This is the current week. In this week, I am hoping to get done with risk assessment and cyber crisis exercise, and then I have to do resilience—cyber resilience—which will take another week. There is a break, unfortunately. This is a big problem, but we will manage. And then I have 3 more weeks here. So, resilience might take 1 week, but also resilience will probably take 2 weeks. And then we can handle STIX maybe in the next week. If there is a need, I will also do an extra one or two classes to cover Zero Trust. So, that should take care of the content of the syllabus. Also, I have posted, I have put out, the recordings that are being done on YouTube links. So, you can watch—well, you are watching it here live. So, you tell your friends that they can watch it from the comfort of their bed. That is the advantage we have—that we have the recordings as well. In any case, I think, let us move on with what we were covering.

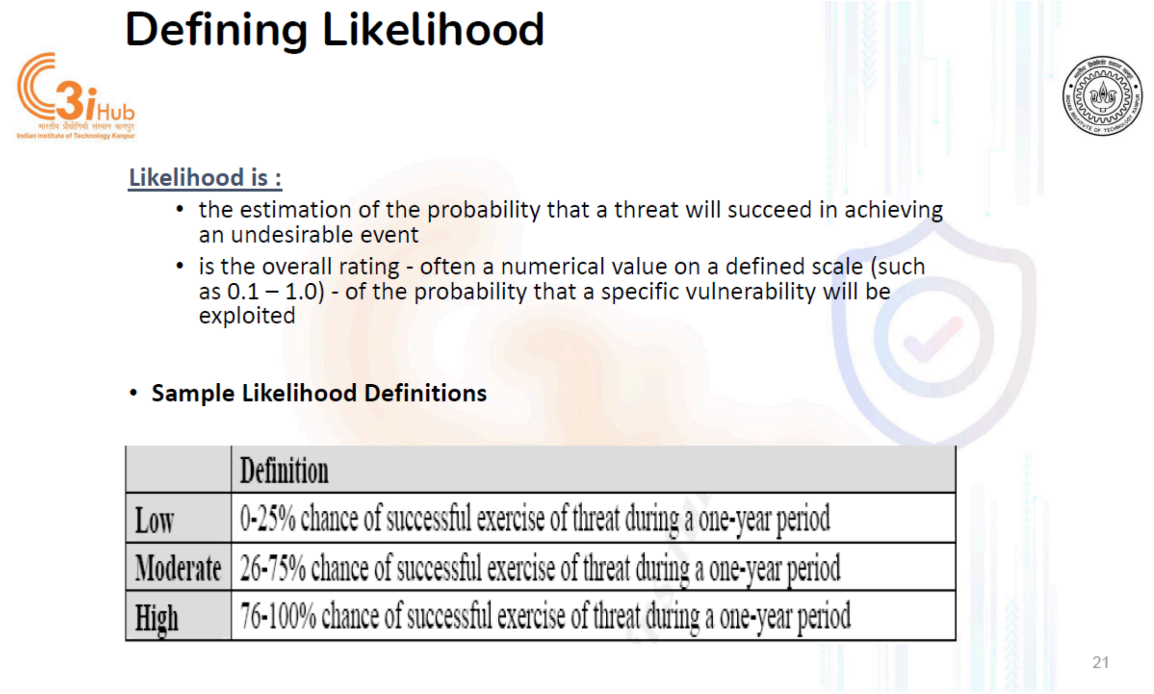

Defining Likelihood

Likelihood is :

- the estimation of the probability that a threat will succeed in achieving an undesirable event
- is the overall rating - often a numerical value on a defined scale (such as 0.1 – 1.0) - of the probability that a specific vulnerability will be exploited

- **Sample Likelihood Definitions**

	Definition
Low	0-25% chance of successful exercise of threat during a one-year period
Moderate	26-75% chance of successful exercise of threat during a one-year period
High	76-100% chance of successful exercise of threat during a one-year period

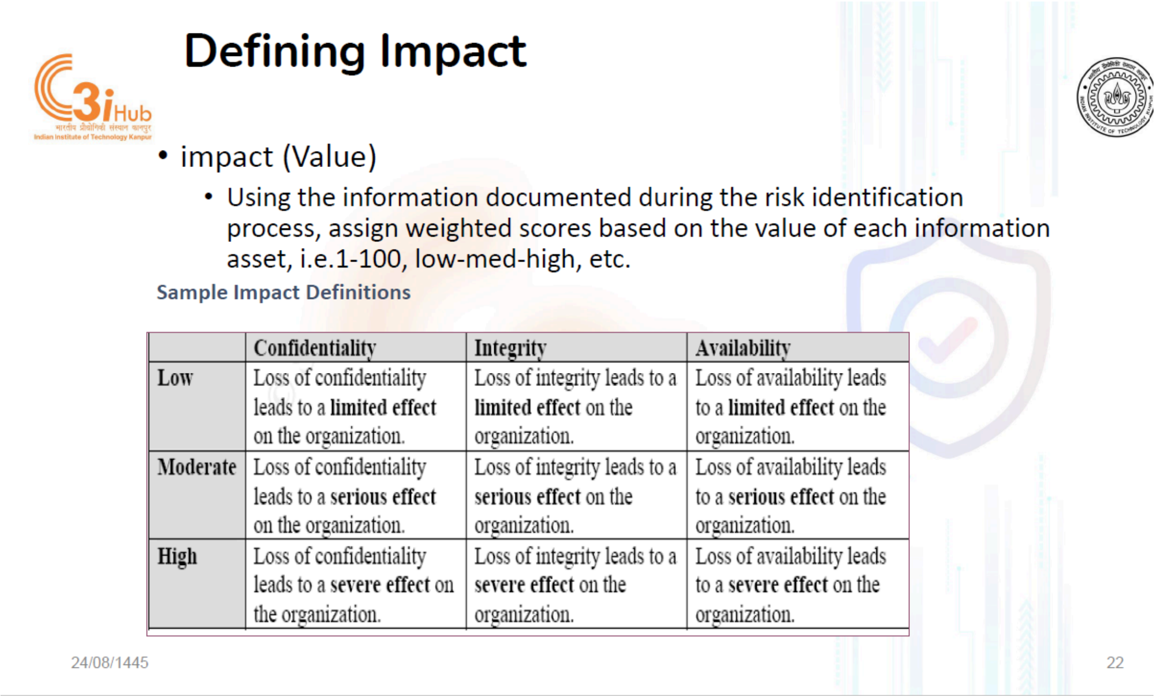

Defining Impact

- impact (Value)
  - Using the information documented during the risk identification process, assign weighted scores based on the value of each information asset, i.e.1-100, low-med-high, etc.

Sample Impact Definitions

	Confidentiality	Integrity	Availability
Low	Loss of confidentiality leads to a **limited effect** on the organization.	Loss of integrity leads to a **limited effect** on the organization.	Loss of availability leads to a **limited effect** on the organization.
Moderate	Loss of confidentiality leads to a **serious effect** on the organization.	Loss of integrity leads to a **serious effect** on the organization.	Loss of availability leads to a **serious effect** on the organization.
High	Loss of confidentiality leads to a **severe effect** on the organization.	Loss of integrity leads to a **severe effect** on the organization.	Loss of availability leads to a **severe effect** on the organization.

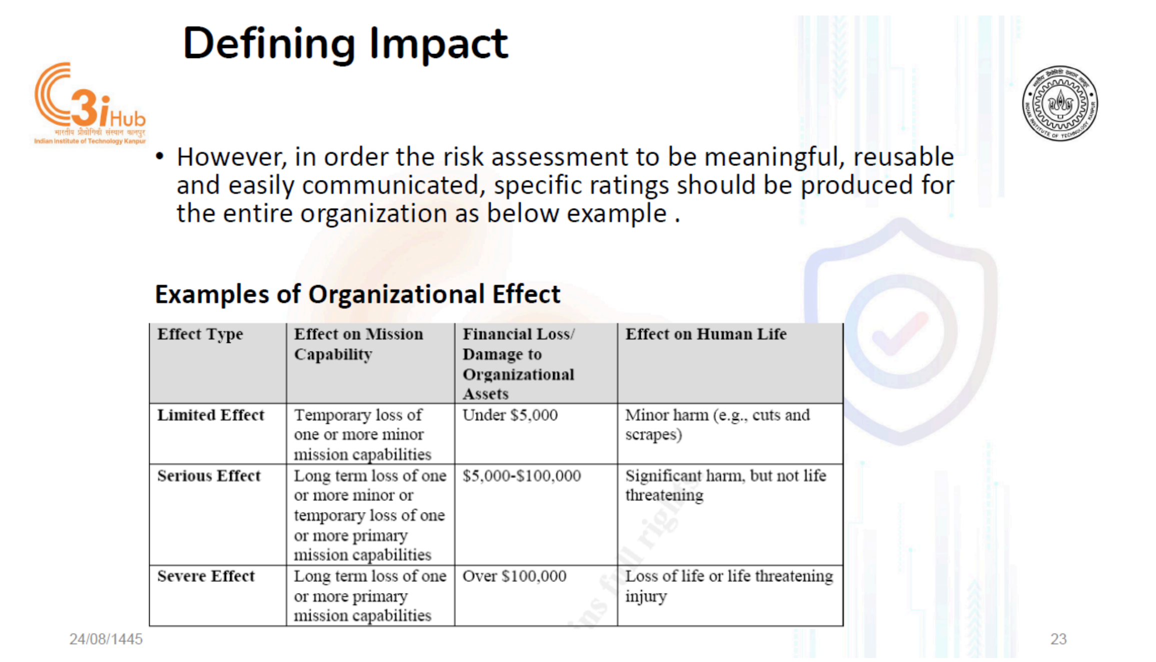

Defining Impact

- However, in order the risk assessment to be meaningful, reusable and easily communicated, specific ratings should be produced for the entire organization as below example .

Examples of Organizational Effect

Effect Type	Effect on Mission Capability	Financial Loss/ Damage to Organizational Assets	Effect on Human Life
Limited Effect	Temporary loss of one or more minor mission capabilities	Under $5,000	Minor harm (e.g., cuts and scrapes)
Serious Effect	Long term loss of one or more minor or temporary loss of one or more primary mission capabilities	$5,000-$100,000	Significant harm, but not life threatening
Severe Effect	Long term loss of one or more primary mission capabilities	Over $100,000	Loss of life or life threatening injury

24/08/1445

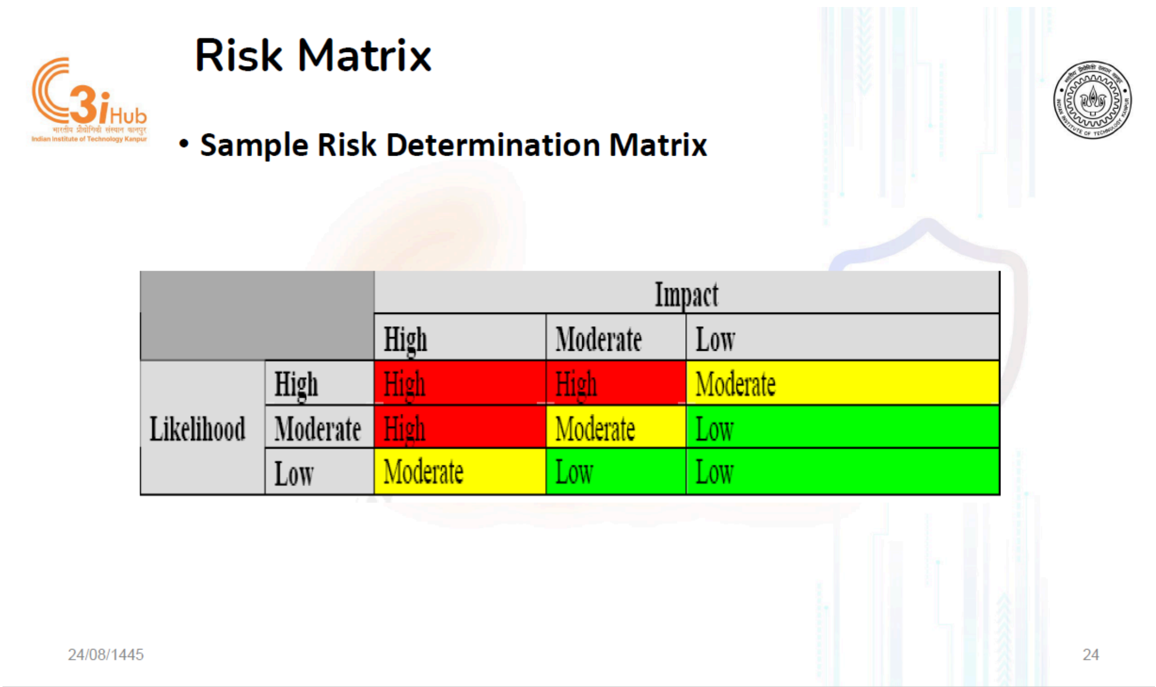

Risk Matrix

- **Sample Risk Determination Matrix**

		Impact		
		High	Moderate	Low
Likelihood	High	High	High	Moderate
	Moderate	High	Moderate	Low
	Low	Moderate	Low	Low

24/08/1445

So, if you remember from last week, we were talking about risk assessment. We talked about how risk assessment requires you to model the threats that are possibly going to harm your business or affect your business processes. Then, we have to figure out whether your system has vulnerabilities such that those threats can actually be realized.

Then, we have to figure out if the threat gets realized, what would be the impact or consequence.

Once you have those three—threats, vulnerabilities—you can estimate the likelihood of a particular threat being realized, and then you find the impact based on the importance of that particular threat, asset, or set of assets that would be affected by that particular threat being realized. What would be the impact? Reputational, financial, business—you know, health and safety of humans and environmental—all kinds of effects might be there. So, that would be your process for risk assessment.

And we also saw that this is done using various metrics, or what we call risk metrics. Also, most of the time, for cyber risk assessment, we do this in a very qualitative manner rather than a quantitative manner. We kind of estimate it as low, high, low, moderate, medium—sorry—low, medium, high, etc. Then, we decide what would be the combination of likelihood and impact strength and finally decide the risk rating of various assets.

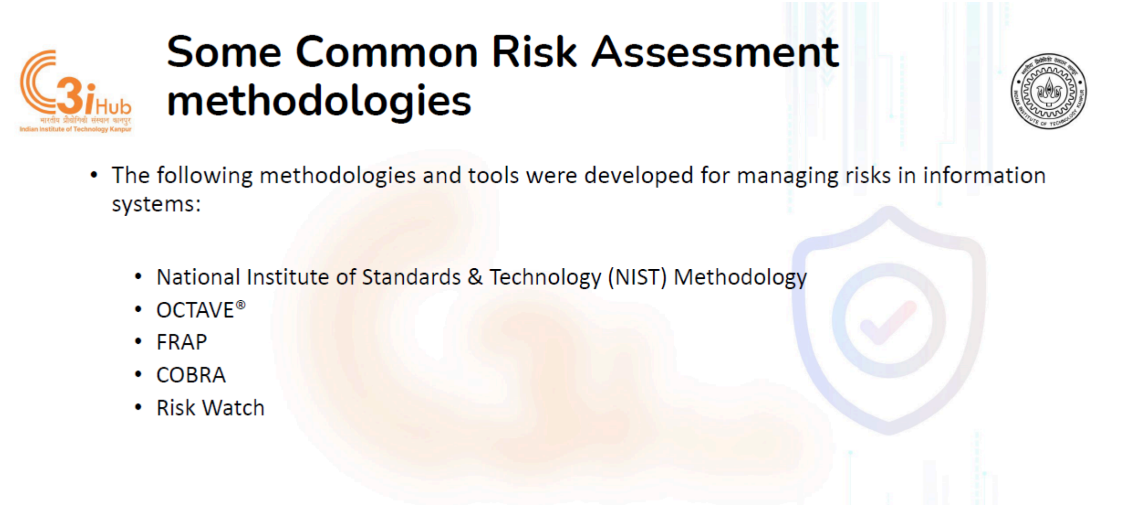

Some Common Risk Assessment methodologies

- The following methodologies and tools were developed for managing risks in information systems:
    - National Institute of Standards & Technology (NIST) Methodology
    - OCTAVE®
    - FRAP
    - COBRA
    - Risk Watch

Now, how to do risk assessment—cyber risk assessment? There are a number of different methodologies. One is the NIST methodology. When we started talking about the NIST methodology, you can also find this document easily on the internet, SP 800-30, which is the NIST Risk Management Guide for Information Technology Systems. And we talked about the steps: you have to understand the system's assets and what assets are there, what their characteristics are. For example, what operating system they are running, what patch level they have, what applications are running, how these assets are connected to each other, whether they are all in the same network, or whether the network has been segmented—all kinds of stuff. Then, you have to figure out what threats are incident on your system or possibly incident on your system, and then vulnerability identification. You have to see whether you have enough mitigating controls, and if you do not, or whatever mitigating controls you have, considering those, you have to figure out what is the likelihood of a particular threat being exercised on your system. Then, you have to figure out what the impact would be if that threat actually affects a particular asset or set of assets. From there, you do risk assessment.

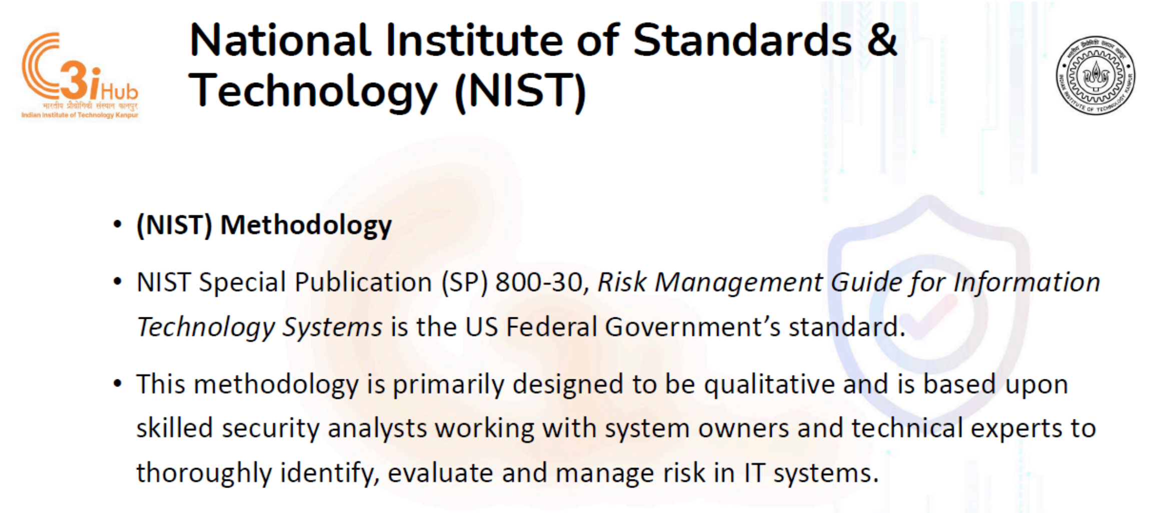

National Institute of Standards & Technology (NIST)

- **(NIST) Methodology**
- NIST Special Publication (SP) 800-30, *Risk Management Guide for Information Technology Systems* is the US Federal Government's standard.
- This methodology is primarily designed to be qualitative and is based upon skilled security analysts working with system owners and technical experts to thoroughly identify, evaluate and manage risk in IT systems.

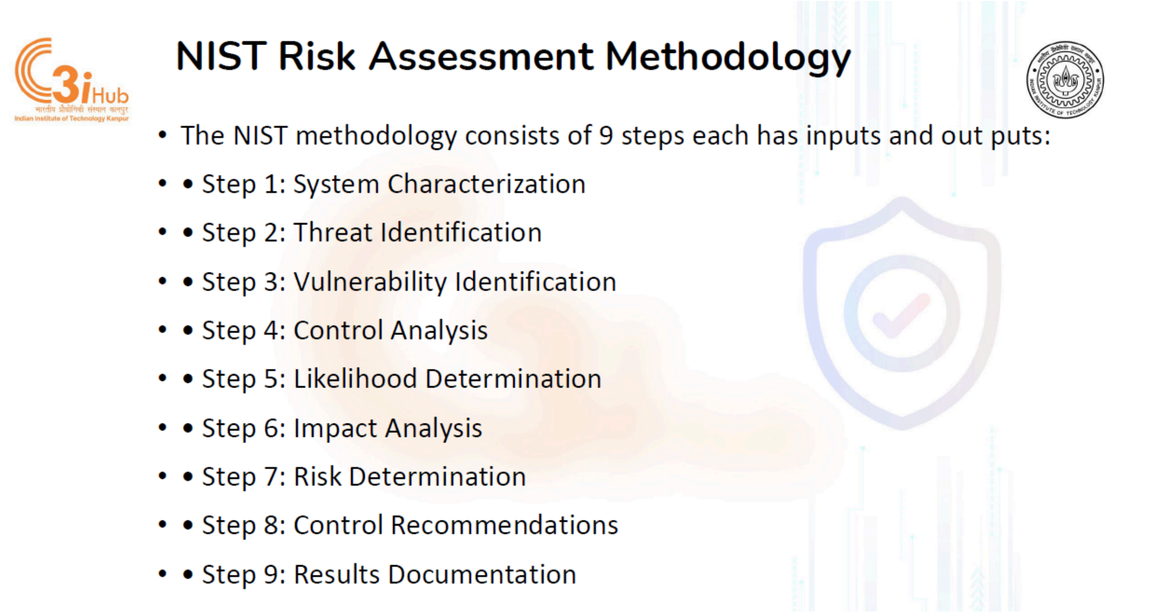

NIST Risk Assessment Methodology

- The NIST methodology consists of 9 steps each has inputs and out puts:
- • 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

Then, you have to say if this risk is acceptable to the organization. If not, then we have to decide what controls need to be there to actually mitigate the risk, to reduce the overall risk to every asset to its acceptable risk level. And then, of course, you have to provide documentation. These are the steps, you know—the inputs and outputs of each step. So, we discuss that here. In the system characterization, you are required to basically collect the asset inventory. So, hardware, software, various applications, system interfaces—people are also assets, and all that. Then, you have to see what the system boundaries are, what the system functions are, the criticality, etc.

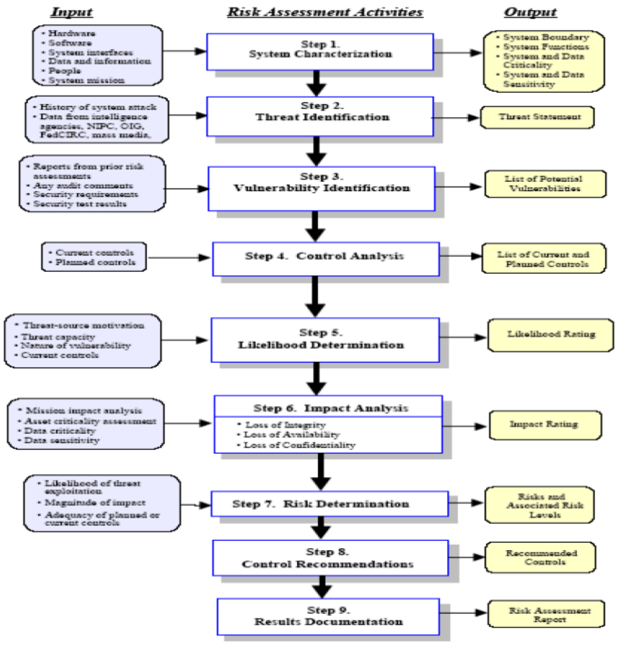

Input | ***Risk Assessment Activities*** | ***Output***

- Hardware
- Software
- System interfaces
- Data and information
- People
- System mission

→ **Step 1. System Characterization** →

- System Boundary
- System Functions
- System and Data Criticality
- System and Data Sensitivity

- History of system attack
- Data from intelligence agencies, NIPC, OIG, FedCIRC, mass media,

→ **Step 2. Threat Identification** → Threat Statement

- Reports from prior risk assessments
- Any audit comments
- Security requirements
- Security test results

→ **Step 3. Vulnerability Identification** → List of Potential Vulnerabilities

- Current controls
- Planned controls

→ **Step 4. Control Analysis** → List of Current and Planned Controls

- Threat-source motivation
- Threat capacity
- Nature of vulnerability
- Current controls

→ **Step 5. Likelihood Determination** → Likelihood Rating

- Mission impact analysis
- Asset criticality assessment
- Data criticality
- Data sensitivity

→ **Step 6. Impact Analysis**
- Loss of Integrity
- Loss of Availability
- Loss of Confidentiality

→ Impact Rating

- Likelihood of threat exploitation
- Magnitude of impact
- Adequacy of planned or current controls

→ **Step 7. Risk Determination** → Risks and Associated Risk Levels

Step 8. Control Recommendations → Recommended Controls

Step 9. Results Documentation → Risk Assessment Report

24/08/1445

28

So, you see what data is there. If that is personally identifiable data, then it has very high criticality already. All this characterization—if a particular server is your database server, which has your critical data, then that server itself becomes critical. If a particular server is actually providing the functionality required by your organization's business, and if that server goes down, and you do not have a failover server, then you will have a problem: your business will be interrupted by a crash of that server. So, you have to declare that server as critical. On the other hand, somebody's desktop—some engineer's desktop—may or may not be critical, depending on what they are doing.

So, all this information needs to be characterized. Then, you have to figure out what threats are possible. For example, you have to know the history of threats, the history of attacks, what kind of attacks have happened. You get threat intelligence from threat intelligence organizations. If you are in a certain type of business, you need to know what threats are happening currently. For example, if you are a bank, there might be a lot of ransomware happening in banks, a lot of unauthorized access to accounts. If you are not too imaginative, these things will help you imagine what threats are possible. If you are imaginative, you could be thinking of not only the threats that are possible from the threat intelligence sources you get but also additional threats that have not happened yet.

So, you can also think of such things, and then you get the threat statements. Then, for vulnerability identification, you have to basically have reports from auditors or VAPT reports, as well as use the national vulnerability database and match it against your applications, operating systems, firmware, etc., for which vulnerabilities are known. Then, you get the list of vulnerabilities.

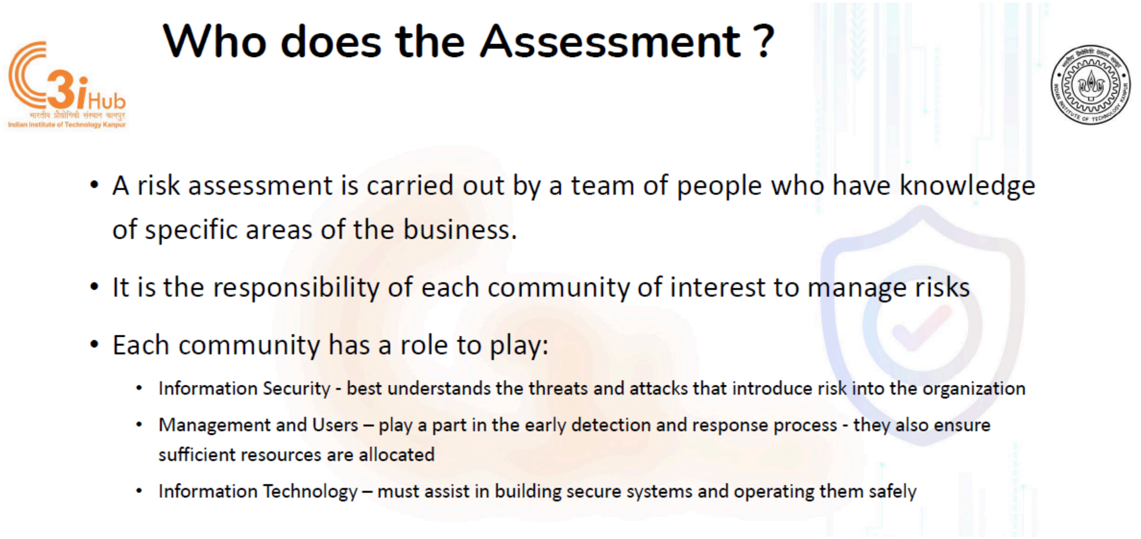

Who does the Assessment ?

- A risk assessment is carried out by a team of people who have knowledge of specific areas of the business.
- It is the responsibility of each community of interest to manage risks
- Each community has a role to play:
  - Information Security - best understands the threats and attacks that introduce risk into the organization
  - Management and Users – play a part in the early detection and response process - they also ensure sufficient resources are allocated
  - Information Technology – must assist in building secure systems and operating them safely

For control analysis, you will actually get the current controls, like where the firewall is, where you have IDS, whether the network is segmented, whether there is endpoint protection and detection on every endpoint, whether there is network intrusion detection, whether there is anti-malware, and whether there are strong authentication and authorization practices. Whether the people have been properly trained not to fall victim to phishing—all these things will be part of your control. So, you say, 'Okay, these are my controls.' In an organization where there has never been any training of personnel on cyber security about phishing, the likelihood that some people will fall victim to phishing and get infected has to be high. Whereas, in an organization that repeatedly does phishing drills and gives awareness training, etc., they might put the likelihood of a phishing-based compromise to be low or medium, right?

So, that is the kind of decision that has to be made in the control analysis. Then, you get the list of current controls. You can also decide that, like, 'Tomorrow, I am going to get this new... well, not tomorrow, but next week, we are going to get this new firewall.' So, it may also be considered a mitigating factor to affect the likelihood. At this point, you have the threats, the vulnerabilities, and the controls, and then you can estimate the likelihood of all the threat-vulnerability pairs, right? Not every threat-vulnerability pair is even possible. So, the likelihood would be 0.

Then, you have to do the impact analysis, like what is the loss incurred by a compromise of a particular asset. This impact analysis will also determine which assets are in the critical business path and which are not. Certain servers, certain database servers, and certain network equipment may be in the critical business path. If one of them gets affected, maybe the entire business process will be interrupted or fail. In such cases, the impact will be considered high, but if it's something on the side, not necessarily involved in the critical business process, then the impact will be low.

Then, you can use the risk matrix to do the risk determination. Here, you have the likelihood estimation, the statements of impact, the controls, and then you can do the risk determination. Then you say, 'Here is your risk: these are high-risk assets, these are medium-risk, and these are low-risk.' Then, you tell the organization, 'Is it acceptable to you?' The organization has a particular preference or decision already made by its management as to what is the acceptable risk. If you can quantify the risk in terms of dollars, that is easier for the management to understand. Like, if this happens, the loss can go into tens of millions, versus if this happens, the loss can go into hundreds of thousands. Maybe you will say, '$100,000, I can accept because the mitigating control you're suggesting will cost me more.' So, they would rather wait for something like this to happen, which is not a good idea, but people think like that. Then, you have to do the final determination.

So, this is what a risk assessor will go into an organization and do. This kind of stuff, right? Risk assessment requires various stakeholders. You cannot get answers to all questions by just talking to the CISO or just one IT head. It requires the information security people, like the CISO's team, the management, and the users. Then, it requires the information technology people who will actually implement the mitigating controls and reduce the risk, right? They will implement the recommendations that you make. So, risk assessment is an important component. Nowadays, all regulators, including the insurance regulator, the stock market regulator SEBI, and the banking regulator RBI, require their regulated entities to do risk assessment, right?

Earlier, what used to happen is that regulators would say you have to implement these controls: you have to have a firewall, strong authentication, two-factor authentication, network segmentation, encryption for data at rest, encryption for data in transit, and antivirus on every endpoint. All these things used to be listed as regulations—cybersecurity regulations. But now, it has been understood that if you are a bank, for example, with tens of thousands of endpoints, multiple branches, hundreds of networks, databases, servers, development machines, and all kinds of connected systems, if you ask them to implement all these things without doing a risk assessment, they will spread their budget across everything.

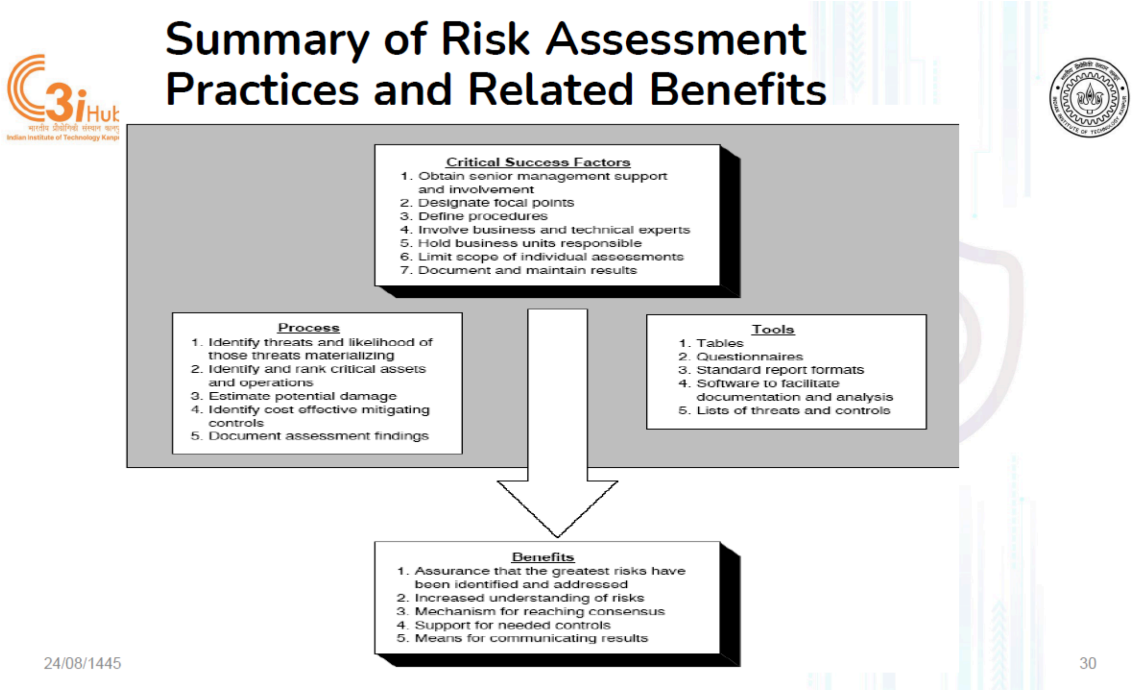

Summary of Risk Assessment Practices and Related Benefits

3iHub
भारतीय प्रौद्योगिकी संस्थान कानपुर
Indian Institute of Technology Kanpur

Critical Success Factors

1. Obtain senior management support and involvement
2. Designate focal points
3. Define procedures
4. Involve business and technical experts
5. Hold business units responsible
6. Limit scope of individual assessments
7. Document and maintain results

Process

1. Identify threats and likelihood of those threats materializing
2. Identify and rank critical assets and operations
3. Estimate potential damage
4. Identify cost effective mitigating controls
5. Document assessment findings

Tools

1. Tables
2. Questionnaires
3. Standard report formats
4. Software to facilitate documentation and analysis
5. Lists of threats and controls

Benefits

1. Assurance that the greatest risks have been identified and addressed
2. Increased understanding of risks
3. Mechanism for reaching consensus
4. Support for needed controls
5. Means for communicating results

24/08/1445

30

But what you want to do is allocate a bigger proportion of your budget to assets that are very critical, that are cyber critical, whereas you can do less for non-critical ones. So, it is very important to actually do the risk assessment, not blindly apply uniform cybersecurity across all organizational assets.

This is again from the NIST risk assessment document. Some of the critical success factors for good risk assessment are actually, first, you have to involve senior management. This is very important, and maybe you will not appreciate this at this point because you might not have worked in a real industry or corporate environment, but getting senior management involved is very difficult. Especially in something like cybersecurity, where many attacks are happening nowadays, top management is getting nervous. They are always worried that something like a ransomware attack may happen, or something like personally identifiable information (PII) may be stolen from customers. So, they are getting more and more involved. It is critical that they get involved, not only because they have to first agree to do a risk assessment. For example, you cannot get an organization like IIT's top management to agree to do a risk assessment because sometimes, we do not want to know that we have a disease.

It's like my mother doesn't want to go to the doctor because she's worried the doctor will say, 'You have this disease,' and then she'll have to take medicine and all that stuff. The same thing happens in organizations where unless the regulator requires them to do a risk assessment, they do not do one. So, top management has to agree to go through a risk assessment because it involves a lot of stakeholders, time, and budget. They also have to designate focal points, like who will be involved, who will be responsible for what, and define the procedures. If you do not get approval from top management, then if you go to a particular sub-organization and ask them for all the asset data and what kind of past attacks they have seen, nobody is going to tell you, right?

So, you need to have a top-down requirement to fulfill in order to get that information. Then, getting business and technical experts is important. Business experts are required because they can only tell you whether this particular asset or business process is crucial for the company. A company may have many business processes, but not all are as crucial as others in terms of funds, income, etc. Holding business units responsible, limiting the scope of individual assessments so that it doesn't delve too deep into unnecessary things, and documenting and maintaining the results from one risk assessment cycle to the next is very important.

The tools that are used are mostly tabular, but there are often questionnaires used to collect data, report formats, software to facilitate documentation and analysis, and lists of threats and controls. There are a number of risk assessment tools available in the market nowadays.

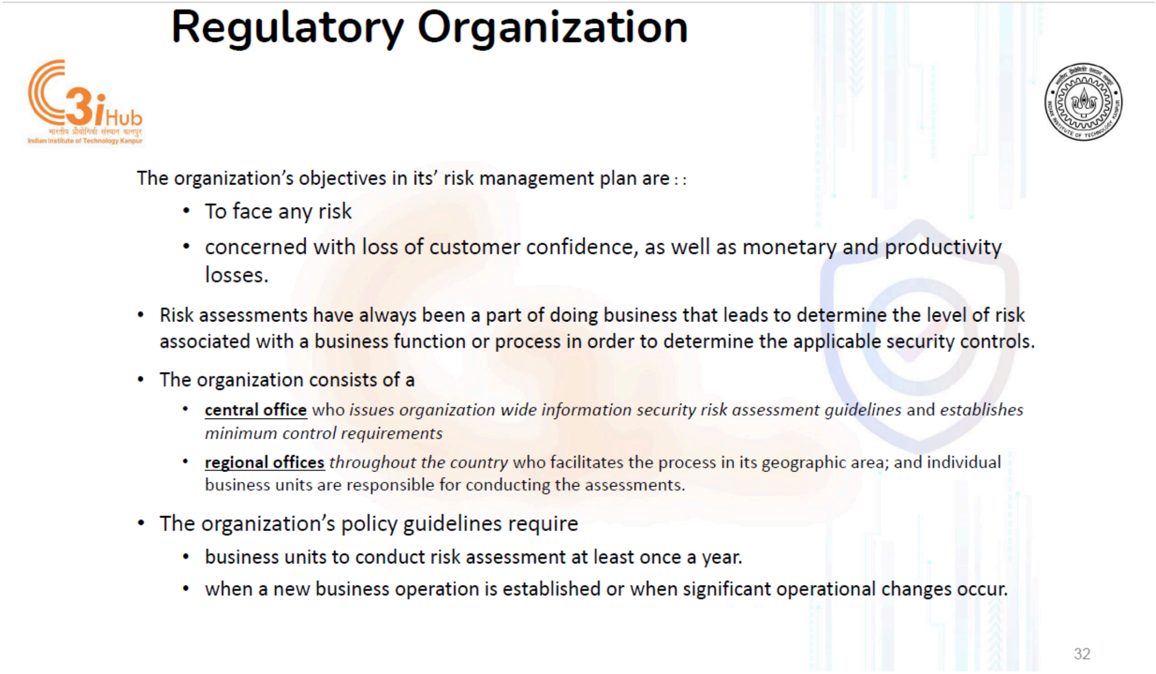

Regulatory Organization

The organization's objectives in its' risk management plan are : :

- To face any risk
- concerned with loss of customer confidence, as well as monetary and productivity losses.

- Risk assessments have always been a part of doing business that leads to determine the level of risk associated with a business function or process in order to determine the applicable security controls.
- The organization consists of a
  - **central office** who *issues organization wide information security risk assessment guidelines* and *establishes minimum control requirements*
  - **regional offices** *throughout the country* who facilitates the process in its geographic area; and individual business units are responsible for conducting the assessments.
- The organization's policy guidelines require
  - business units to conduct risk assessment at least once a year.
  - when a new business operation is established or when significant operational changes occur.

32

There are some startups that we, C3i hub, support, which have built or are building risk assessment tools, right. Similarly, there are companies like Risk Lens, which has now been bought by Safe Security. There is an organization called Fair Institute, which is a risk assessment methodology school of thought. So, there are lots of different tools nowadays for helping the risk assessors to gather information. Eventually, it is a lot of manual work. It cannot be fully automated, where I plug in an agent on every device on the system, and it gives me data and a risk indicator, but I cannot get a full-fledged risk assessment through that.

There is also a set of processes, as you have seen, like identifying threats, identifying vulnerabilities, identifying the criticality of assets based on their importance in the business process, the cost of mitigating controls, all that stuff.

Of course, there is a benefit of risk assessment. We have been saying why we need to do risk assessment; we need to understand where to actually put our security budget. So, that is what it is. Now, let us quickly look through some case studies. These case studies are not very elaborate, but they will give you some idea. So, here is a regulatory organization that is trying to do its risk assessment.

The goal for doing the risk assessment is actually to face any risk, and they are concerned with the loss of customer confidence, monetary loss, and productivity losses. As I said before, the term 'risk' is often overloaded, right. So, the term 'risk' when we say risk assessment versus when we use the term 'risk' in an everyday sense is quite different. For example, if we are driving, we say that we have a risk of an accident, right. But actually, risk is supposed to be a measure of that particular threat being realized and the impact, right. So, the risk, if you are the sole breadwinner of your family, is higher, right, whereas if you are not.

Now, risk is not higher. This is a very controversial statement because, for whoever you are, for your parents and your children or wife or husband, you are very important. So, I should not estimate that in those terms. In terms of the economy, that is the case. But in any case, here, when they say that they want to face any risk, what they mean is they want to face any threat, not risk, right. Because these are threats that they will lose customer confidence, or they will have monetary loss, or they will have productivity loss, etc. These are actually threats.

Now, risk assessment—so, this is an organization that has been doing risk assessment as a part of their regular business, and the level of risk associated with the various business processes is actually a regular part of assessment. Now, the organization has a central office and a regional office, and the central office decides what the requirements, processes, and rules for risk assessment are for the entire organization, and regional offices actually get it done for all the business units within the regional office. The organization’s policy guidelines say that the business units have to conduct risk assessment at least once a year or when a new business operation or some kind of significant operational change happens, right. So, this is a normal rule for most organizations, that either you do it periodically, but in case you are going to have a significant change in your organization—for example, you were only doing brick-and-mortar business, and you decide to also do e-commerce business—then you want to do a risk assessment at that point in time.

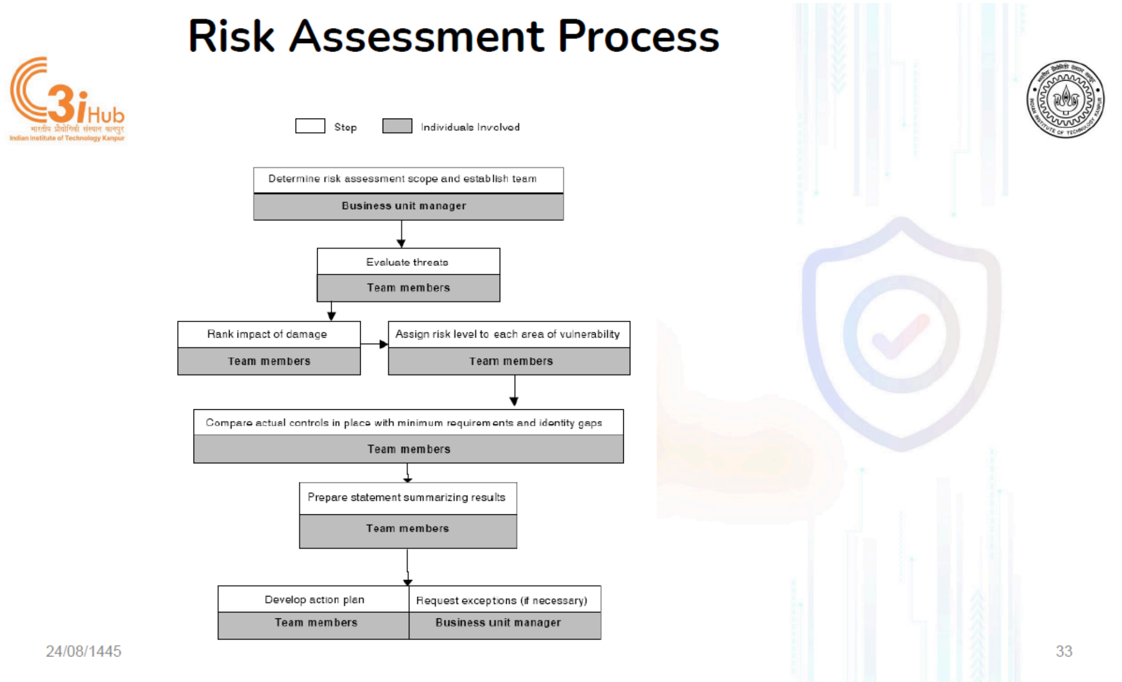

Risk Assessment Process

Determine risk assessment scope and establish team — Business unit manager

Evaluate threats — Team members

Rank impact of damage — Team members

Assign risk level to each area of vulnerability — Team members

Compare actual controls in place with minimum requirements and identify gaps — Team members

Prepare statement summarizing results — Team members

Develop action plan — Team members

Request exceptions (if necessary) — Business unit manager

24/08/1445

33

So, that is the periodicity. And here is this picture that shows you what the different activities are and who is responsible.

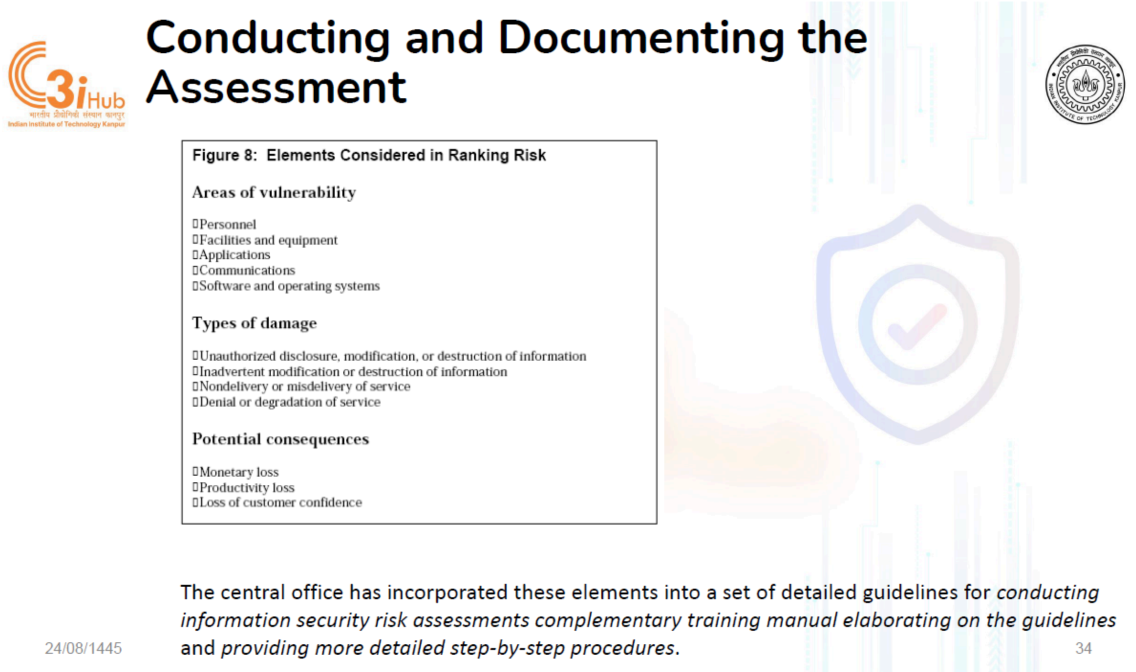

Conducting and Documenting the Assessment

Figure 8: Elements Considered in Ranking Risk

Areas of vulnerability

- Personnel
- Facilities and equipment
- Applications
- Communications
- Software and operating systems

Types of damage

- Unauthorized disclosure, modification, or destruction of information
- Inadvertent modification or destruction of information
- Nondelivery or misdelivery of service
- Denial or degradation of service

Potential consequences

- Monetary loss
- Productivity loss
- Loss of customer confidence

The central office has incorporated these elements into a set of detailed guidelines for *conducting information security risk assessments complementary training manual elaborating on the guidelines* and *providing more detailed step-by-step procedures.*

24/08/1445

34

As you can see, the business unit manager decides the assessment scope and who is responsible for the assessment, and then you have the rest of the process, like evaluation of threats, evaluation of vulnerabilities, figuring out the impact, etc., all done by the team members. There is nothing interesting to see in this diagram. And then here is what the organization decided as its areas of vulnerabilities and the kinds of damages that can happen.

Here, you can see that areas of vulnerabilities are, like, personnel, which were probably the biggest concern in many organizations. Then you have the various facilities, and you have various applications, communication, and eventually the software and operating systems. And the type of damage you can see is the unauthorized disclosure, right. So, they are more of an information organization, right.

So, they probably do not sell any products or something. So, they have unauthorized information disclosure and/or modification. Unauthorized disclosure is a confidentiality breach; modification is an integrity breach. And then you have the destruction of information, so that is an availability breach.

You have inadvertent disclosure and modification. The difference between the two is that one is a willful, malicious breach of confidentiality, integrity, etc., and the other is inadvertent, which can happen if you do not have the right control. For example, if you do not have backups, then you can have inadvertent destruction of information. If you do not have hashes or some other way to actually have integrity checks, then you might have inadvertent modification. If you do not keep the data in encrypted form, then you may have inadvertent disclosure, right. So, those are the kinds of things that can happen. Non-delivery or misdelivery of service can happen.

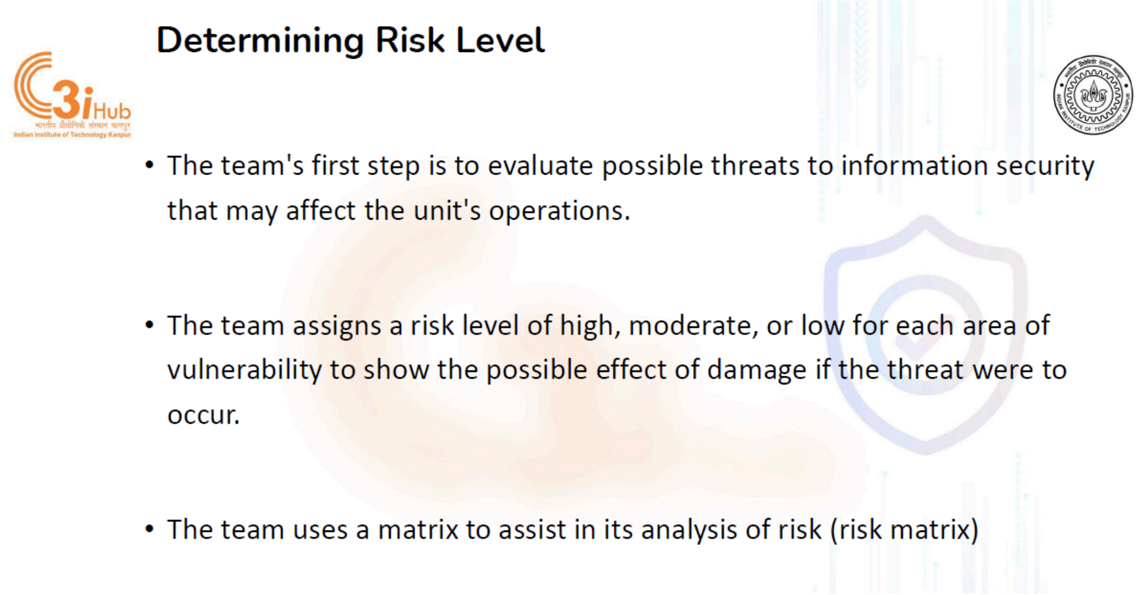

Determining Risk Level

- The team's first step is to evaluate possible threats to information security that may affect the unit's operations.
- The team assigns a risk level of high, moderate, or low for each area of vulnerability to show the possible effect of damage if the threat were to occur.
- The team uses a matrix to assist in its analysis of risk (risk matrix)

So, either because of malicious activities or because something went wrong, you can have denial or degradation of service. So, these are the damages that can happen either by employees being targeted or by virtue of some software, application, communication channels, or facility being targeted. This is a very high-level risk assessment, not at the level where we actually go for CVEs and specific threat scenarios. So, this will give you a very high-level idea about the risks. And the consequences are basically monetary loss, production loss, and loss of customer confidence.

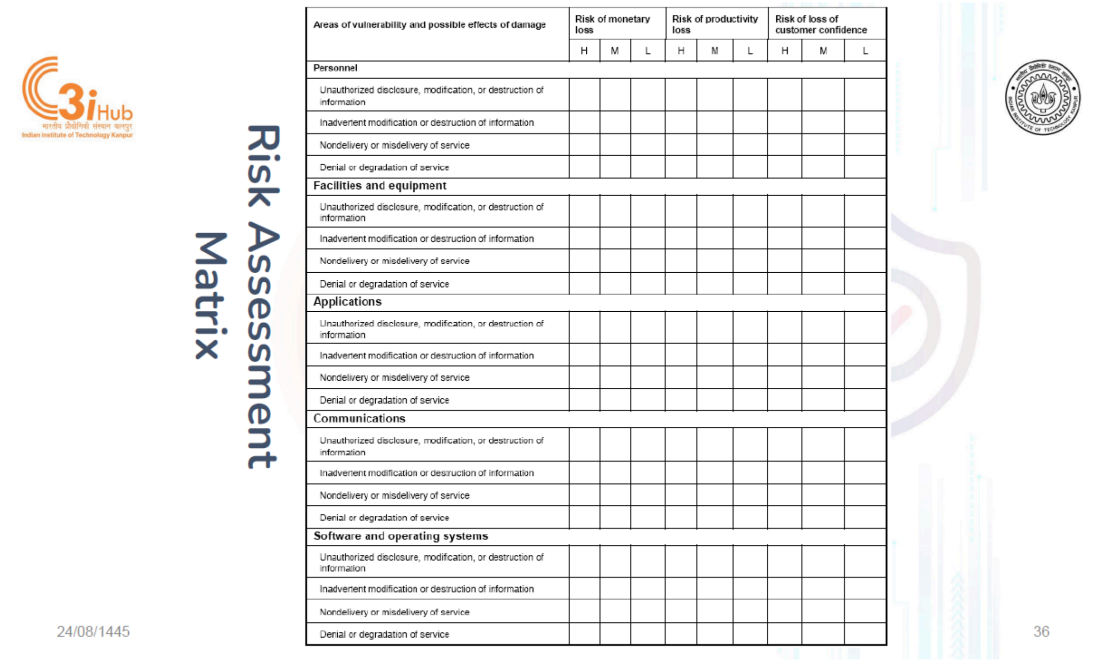

Risk Assessment Matrix

Areas of vulnerability and possible effects of damage	Risk of monetary loss			Risk of productivity loss			Risk of loss of customer confidence		
	H	M	L	H	M	L	H	M	L
Personnel									
Unauthorized disclosure, modification, or destruction of information									
Inadvertent modification or destruction of information									
Nondelivery or misdelivery of service									
Denial or degradation of service									
Facilities and equipment									
Unauthorized disclosure, modification, or destruction of information									
Inadvertent modification or destruction of information									
Nondelivery or misdelivery of service									
Denial or degradation of service									
Applications									
Unauthorized disclosure, modification, or destruction of information									
Inadvertent modification or destruction of information									
Nondelivery or misdelivery of service									
Denial or degradation of service									
Communications									
Unauthorized disclosure, modification, or destruction of information									
Inadvertent modification or destruction of information									
Nondelivery or misdelivery of service									
Denial or degradation of service									
Software and operating systems									
Unauthorized disclosure, modification, or destruction of information									
Inadvertent modification or destruction of information									
Nondelivery or misdelivery of service									
Denial or degradation of service									

24/08/1445

36

So, these are the consequences. Now, the question is how do you determine the risk level. We will go through this matrix method. This is a very high-level risk assessment. As you can see, on the left-hand side (on the rows), you have the various vulnerability areas, like people, application, software, facilities, etc., and for each of these, you have various threats.

So, it is a threat-vulnerability pair, right. The threat-vulnerability pair is on the rows, and on the columns, you have the likelihoods. The likelihoods are high, medium, low. But instead of computing the likelihood altogether, they are actually doing it with respect to three different parameters: what is the likelihood of monetary loss, what is the likelihood of production loss, and what is the likelihood of the loss of customer confidence. By filling in this table, they can then summarize it. So, they will summarize, for each area of vulnerability, what the risk is. Because if the monetary loss is low and the productivity loss and customer confidence loss are high, then maybe you will put the overall risk as high. So, this is how they are computing the overall risk. Then, you have to find what controls you should have for protection against these losses. And then eventually, somebody has to take action, and this action will then be checked by the same team that did the risk assessment to see whether the risk has come down.

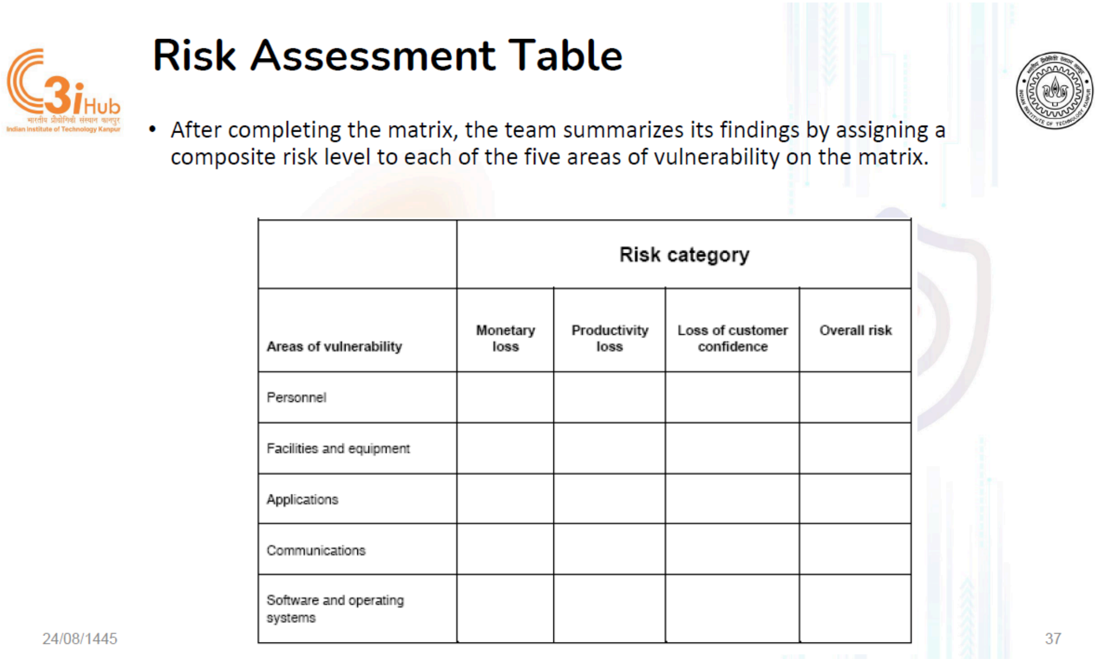

Risk Assessment Table

- After completing the matrix, the team summarizes its findings by assigning a composite risk level to each of the five areas of vulnerability on the matrix.

	Risk category			
Areas of vulnerability	Monetary loss	Productivity loss	Loss of customer confidence	Overall risk
Personnel				
Facilities and equipment				
Applications				
Communications				
Software and operating systems				

24/08/1445

37

Now, from this analysis, can you actually do something actionable? This analysis does not tell you anything about actual organizational IT assets, right. It does not tell you what the network architecture is, right. It does not tell you what kind of information it possesses, right. It does not tell you anything about the patches and versions of the software, firmware, hardware, etc.

Therefore, when we are making decisions like application risks, we are saying, 'Okay, unauthorized disclosure, modification, and destruction of information,' and we are saying, 'Let’s say the corresponding likelihood of monetary loss is high,' right. But we are not really knowing whether this is even possible; maybe that likelihood is 0, right. So, I am just doing it at a very high level, where we are assuming the worst. We are saying, 'Okay, if there is an application that is important for the business process, and if it discloses any information like this, then maybe my monetary loss will be high, customer confidence loss will be high.' But it does not really say what the vulnerability in this organization is that would lead to this kind of disclosure, right.

So, this is what is called a very high-level risk assessment. This only shows what would happen in the worst case, provided there are no controls and all the software, etc., are outdated and have vulnerabilities. So, this kind of analysis you can get from this level. Without having a realistic picture, which is called a detailed risk analysis, you cannot make accurate decisions. A detailed risk analysis will require you to have the asset inventory, the exact version of all the software, hardware, and firmware. You have to get the vulnerabilities from the NVD database; you have to get the vulnerabilities that your last audit report mentioned and that you have not fixed yet. You have to have the full details of the controls you already have, and only then can you do this analysis. That is the detailed analysis. A detailed analysis will give you the nuanced differences between multiple different servers; they will have different criticality.

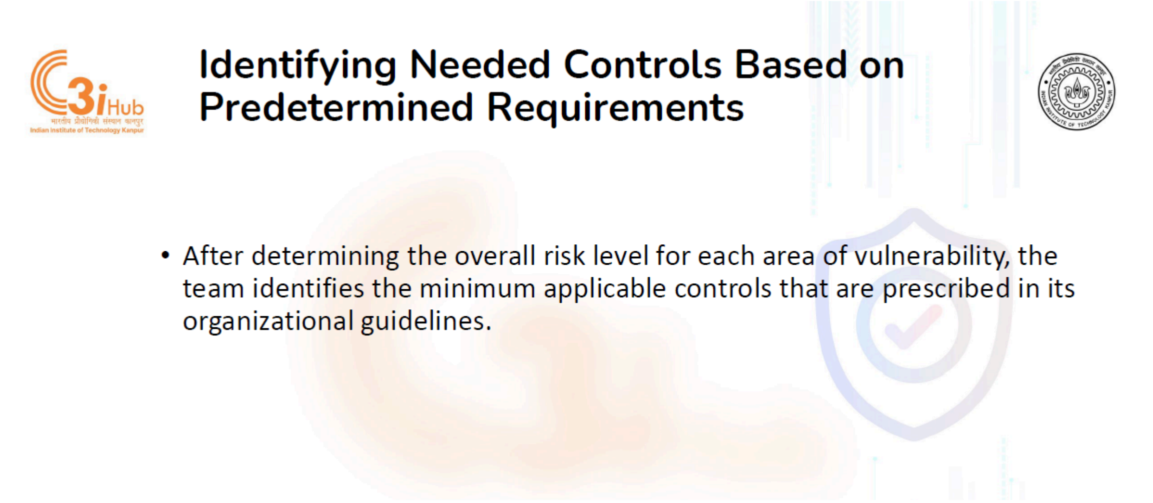

Identifying Needed Controls Based on Predetermined Requirements

- After determining the overall risk level for each area of vulnerability, the team identifies the minimum applicable controls that are prescribed in its organizational guidelines.

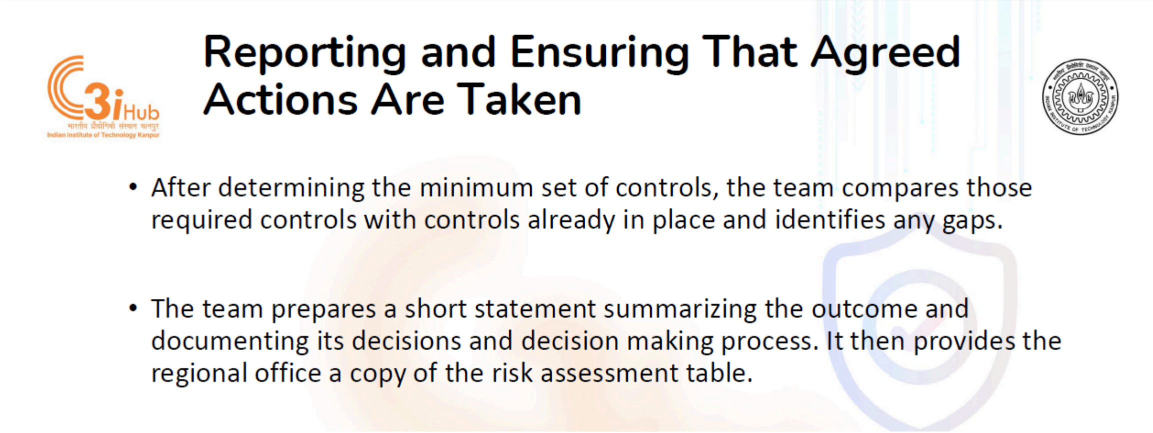

Reporting and Ensuring That Agreed Actions Are Taken

- After determining the minimum set of controls, the team compares those required controls with controls already in place and identifies any gaps.
- The team prepares a short statement summarizing the outcome and documenting its decisions and decision making process. It then provides the regional office a copy of the risk assessment table.

But in this high-level analysis, everything is going to be either critical or non-critical, right. So, this is a very high-level risk assessment. Another one is also a very high level, but here it is actually not even a proper risk assessment; the goal is to have a security plan. So, this is a university that wants to decide whether it has an adequate cybersecurity plan for information, because unlike India, in the US, student information is protected information. In the US, if your parent comes to me and asks for your grade, I cannot tell them without your consent.

Here, I can tell your grade to your parents if they ask, and they do ask sometimes, but in the US, I could not do that without the student’s consent, because students are above 18 in the university, and I cannot say. So, there is a federal law that says student information is very critical and important. Their grades, their social security numbers, roll numbers, addresses—everything is protected. When that law was enacted, I remember we were given special software by the university to scrub our emails because our past emails might have information in which a student’s name and roll number might be in the same email. That was no longer allowed. So, if a student’s name is there, no other information that relates to a database key is allowed, right.

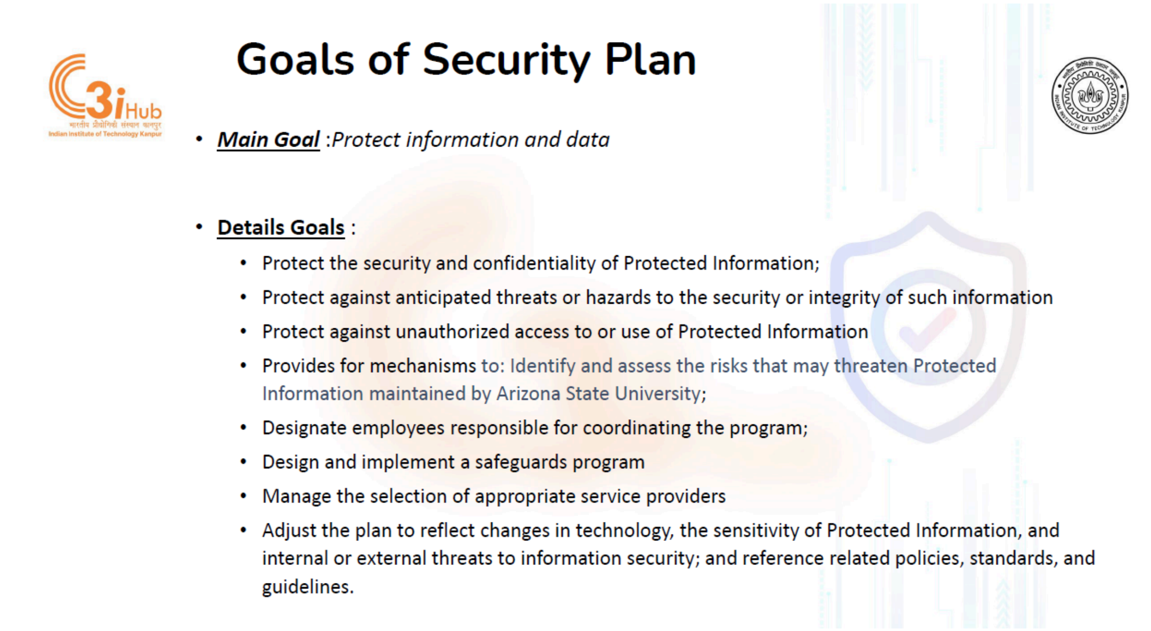

Goals of Security Plan

- ***Main Goal*** :*Protect information and data*

- **Details Goals** :
  - Protect the security and confidentiality of Protected Information;
  - Protect against anticipated threats or hazards to the security or integrity of such information
  - Protect against unauthorized access to or use of Protected Information
  - Provides for mechanisms to: Identify and assess the risks that may threaten Protected Information maintained by Arizona State University;
  - Designate employees responsible for coordinating the program;
  - Design and implement a safeguards program
  - Manage the selection of appropriate service providers
  - Adjust the plan to reflect changes in technology, the sensitivity of Protected Information, and internal or external threats to information security; and reference related policies, standards, and guidelines.

So, we had to scrub all our emails and everything, all the files in our personal desktops, to actually make sure that there is no file in which a student’s name and roll number are in the same document. I can have a document in which the roll number and grade are in the same document, but if the name is also there, then it will become a risk. So, we had to scrub. This is the situation. Under this situation, let us say a university like Arizona State University is trying to actually protect the information.

So, security and confidentiality of the information. They want to protect against anticipated threats or hazards to the security and integrity of the information. Unauthorized access—they do not want. So, they do not want somebody accessing it. There was a case in the US where a student actually put a key logger on machines in the library—all the machines. The library had a lot of machines, and this was some time ago. At that time, people would actually go to a library website, use a library computer, and do some work there.

No longer does anybody do that. It seems professors—this is in Colorado—went to the library and accessed the grade database with the professors' login, and there was no two-factor login. This student, using the key logger, got all the passwords from all the professors. What he did was that he changed all his grades to A, right. Now, his name came in the dean's list. When his name came in the dean's list, all the professors were very surprised because this was a student they thought was an F student, and suddenly he seems to have a 4 out of 4 GPA. Then, they started digging, and they figured out from the logs that somebody went in and just changed that guy's grades in multiple different courses to A.

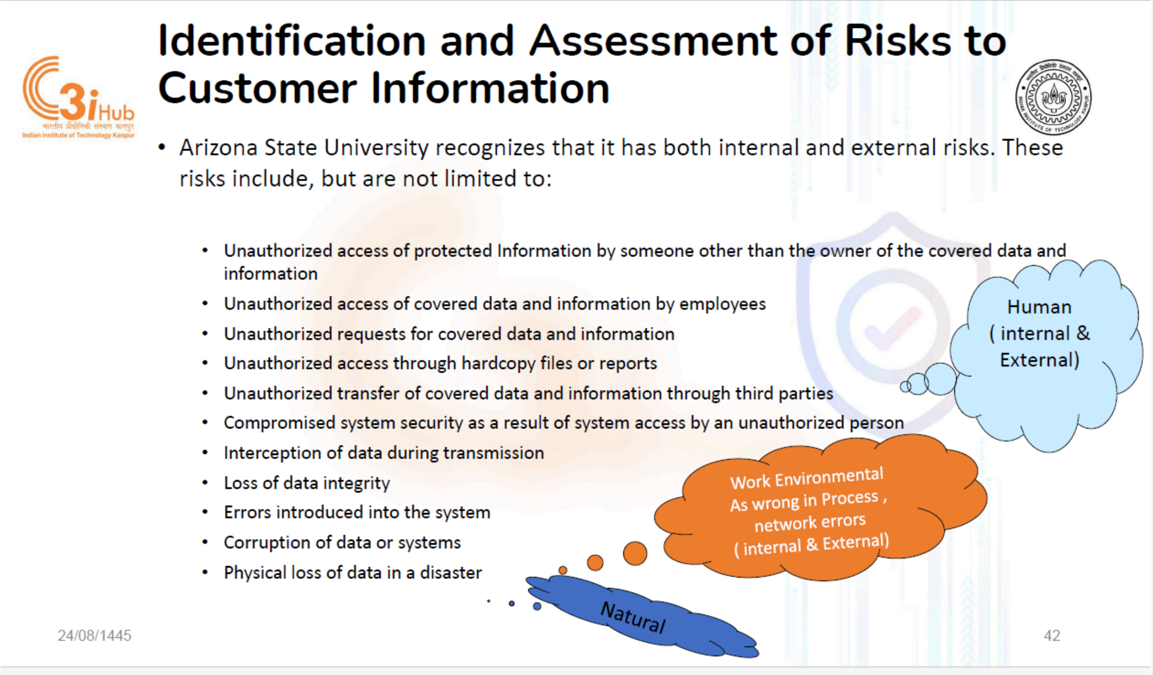

Identification and Assessment of Risks to Customer Information

- Arizona State University recognizes that it has both internal and external risks. These risks include, but are not limited to:
  - Unauthorized access of protected Information by someone other than the owner of the covered data and information
  - Unauthorized access of covered data and information by employees
  - Unauthorized requests for covered data and information
  - Unauthorized access through hardcopy files or reports
  - Unauthorized transfer of covered data and information through third parties
  - Compromised system security as a result of system access by an unauthorized person
  - Interception of data during transmission
  - Loss of data integrity
  - Errors introduced into the system
  - Corruption of data or systems
  - Physical loss of data in a disaster

Human (internal & External)

Work Environmental As wrong in Process , network errors (internal & External)

Natural

24/08/1445

42

Now, if the student was clever, he would have done something clever, like giving himself a C in everything, then he would not have come onto the dean's list, and no professor would be noticing it. But unfortunately, he went a little too aggressive, and therefore he actually got caught, right. So, this is what we do in intrusion detection—we always try to make the intrusion stealthy. The intrusion should go under the radar. It should not be so anomalous that the detection system triggers an alert, right. In any case, the student obviously went to jail and everything. Interestingly, it was an Indian-origin student, right. But anyway, this is what we want to do: unauthorized access—we want to stop it. We want the mechanism to identify and assess the risks for the protection of this information, and we want to designate employees who are responsible for coordinating the security program.

We also want to design and implement a safeguard program, manage the selection of appropriate service providers, like who would be a firewall vendor, who will be my intrusion detection vendor, who would be my SOC and SIEM vendor, antivirus vendor, etc., and adjust the plan to reflect changes in technology. So, the plan has to be dynamically adapted as technology changes, the threat landscape changes, and so on and so forth. So, that is what the organization wants to do. The first thing they have to do is figure out where the risks are. When it says 'What are the internal and external risks?' what they mean, again, is a misuse of the term—it should be 'What are the internal and external threats?' So, you can easily imagine what the internal and external threats are: unauthorized access, inadvertent disclosure, interception of data during transmission, loss of data integrity, errors introduced in the system, corruption of data or systems, physical loss of data in a disaster like a flood, hurricane, tornado, etc.

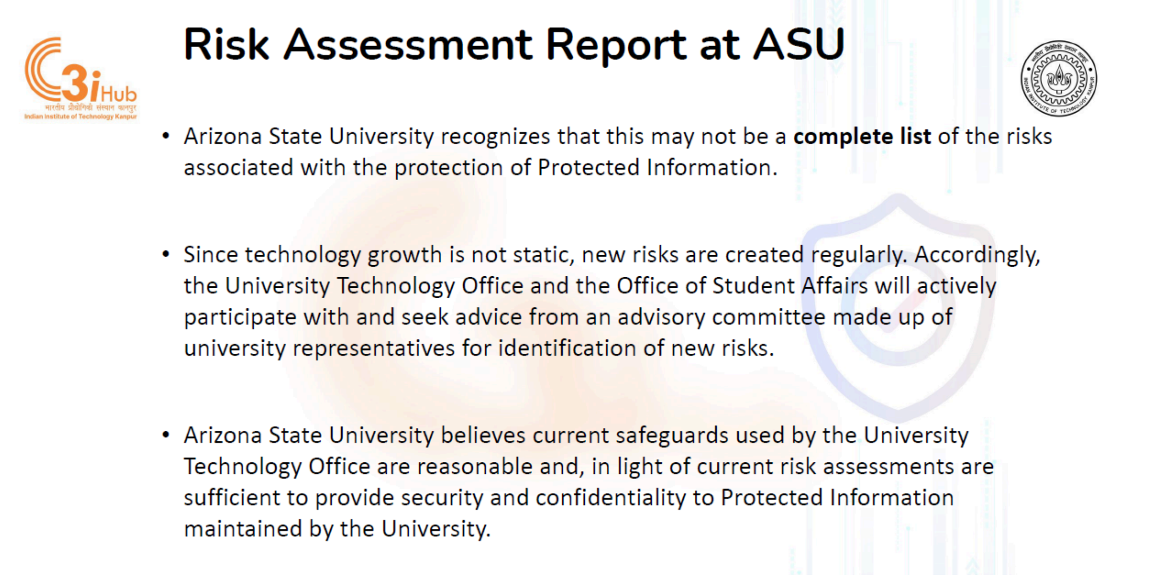

Risk Assessment Report at ASU

- Arizona State University recognizes that this may not be a **complete list** of the risks associated with the protection of Protected Information.
- Since technology growth is not static, new risks are created regularly. Accordingly, the University Technology Office and the Office of Student Affairs will actively participate with and seek advice from an advisory committee made up of university representatives for identification of new risks.
- Arizona State University believes current safeguards used by the University Technology Office are reasonable and, in light of current risk assessments are sufficient to provide security and confidentiality to Protected Information maintained by the University.

So, they commissioned a risk assessment, and they figured out that, okay, a risk assessment may not get a complete list of threats. As I said, some threats you can model or imagine based on threat intelligence reports, threat intelligence feeds, past attacks, etc. But certain things may be totally new, and you might not have thought about them, like the student putting key loggers on the library desktops—that was not thought about before it happened, right. So, you may not have a complete list, but you try to do comprehensive threat modeling, right.

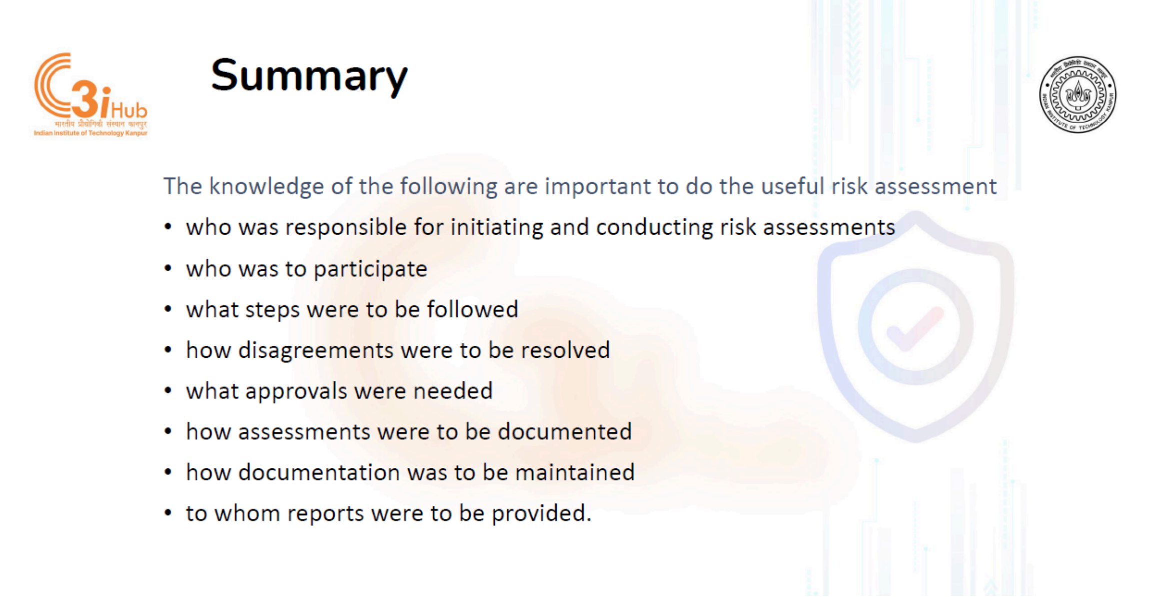

Summary

The knowledge of the following are important to do the useful risk assessment

- who was responsible for initiating and conducting risk assessments
- who was to participate
- what steps were to be followed
- how disagreements were to be resolved
- what approvals were needed
- how assessments were to be documented
- how documentation was to be maintained
- to whom reports were to be provided.

You want to understand the directions from which attacks may come. They figured out that since technology growth is not static, new risks—new threats—are created regularly. Accordingly, the officers of the technology office and the office of student affairs will create an advisory committee of experts, and they will have representatives from various departments. They will always have regular evaluations of risks by considering new threats and new vulnerabilities. Maybe they will have new business processes—for example, the university earlier might not have had online teaching, but they might have added online teaching. That would be a new technology, a new source of threat, and a new source of vulnerability, and therefore, they have to worry about that.

So, who is responsible for assessing the risk? The technology officer, with the advisory committee, will include representatives from all departments. Each department will provide an annual update report about its safeguarding procedures. This is what is called cybersecurity governance, right. Cybersecurity, as I said—well, I don't know if I said it or not, but I should have said it—is about three things: people, process, and technology, right. Technology is one thing that you normally do in terms of intrusion detection, in terms of firewalls, and in terms of endpoint protection and detection, or antivirus. You do this from the SOC, SIEM, and all kinds of things.

But people and process are also equally important, and cybersecurity governance is very, very important. This is all about governance, right. Now, what you will see—maybe next week, not this week—is that the NIST framework, the National Institute of Standards and Technology, has a cybersecurity framework. In the framework, they actually added a new function. They used to have five functions: identify, protect, detect, respond, and recover. Now, they have added a sixth: govern.

Governance is very, very important, and what this is talking about is governance. So, this is where I will stop in terms of risk assessment. You get the basic idea of risk assessment, but you will get a better idea when the homework is released, and you will actually face the issues that a risk assessor will face. Then, you will get a better understanding of risk assessment. So, we will stop the risk assessment here and move on to the cyber crisis exercise in the next class.