Practical Cyber Security for Cyber Security Practitioners

Prof. Sandeep Kumar Shukla

Department of Computer Science and Engineering

Indian Institute of Technology, Kanpur

Lecture 22

All right, so good morning. We have been talking about cyber crisis management exercises, and a cyber crisis management exercise is something that organizations do in order to test their ability to respond to a cyber crisis. As we discussed, a cyber crisis is a rather rare incident, but once it happens, the consequences are very high. It basically disrupts the business processes. So, whatever the business your organization is in, you have to figure out what kind of disruption would constitute a crisis. Then, accordingly, you have to have a crisis management plan, and that management plan has to be well understood by all stakeholders, including the topmost level of management. When a crisis happens, decision-makers must first decide whether the incident they are dealing with is actually a crisis, rather than just an incident. But the bigger issue is that once you declare it is a crisis, certain mechanisms, a certain response playbook, must come into effect. Once that playbook comes into effect, everyone who has been given responsibility to do certain things must do that.

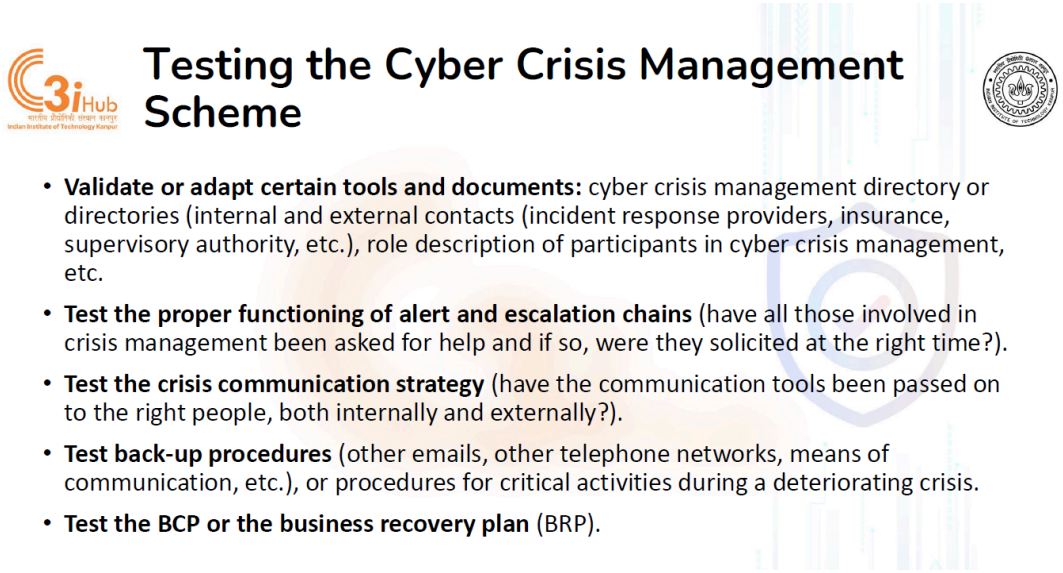

For example, the IT team may need to shut down certain applications, shut down certain computing servers or internet-facing servers, or segment a part of the network that has

been affected, and so on. Similarly, the cybersecurity team may need to collect data for further forensic analysis in order to understand what is actually affected and what is not. The communication team also has to decide whether to report the incident to the regulatory authority. For example, CERT-In requires certain types of incidents to be reported within six hours. If reporting is necessary, the team must collect the right set of information to be communicated to the regulator so the regulator can decide whether to inform other regulated entities in the same sector to be cautious and on guard against this kind of threat. Additionally, they have to figure out how to communicate this to the outside world, including the press. In India, not much is done in terms of informing the press, but, for example, in the US, the Securities and Exchange Commission (the organization equivalent to SEBI, the Securities and Exchange Board of India) requires that any kind of cyberattack affecting business continuity be reported to the SEC, as it might affect the stock market value of the company. Thus, the communication part has to be correct: you need to decide what to communicate, how much to communicate, and to whom. You also need to communicate with the rest of the organization about what to do and what not to do, such as whether to log into their systems or disconnect them, and what kind of actions to take until the crisis has dissipated. This is what a cyber crisis exercise helps with, preparing you to drill and be ready for this scenario.

What we are saying here is that testing cyber crisis management allows us to validate or adapt certain tools and documents. We must also document everything that needs to be done so that everyone knows exactly what to do when something like this happens. The exercise helps to debug the document—whether it contains the right set of principles and steps. Without the exercise, you would only know when a real crisis happens, which could pose a problem.

They also need to test the proper functioning of alerts and the escalation chain. Typically, there will be a SOC (Security Operations Center) where SOC analysts are present 24/7, monitoring alerts on the screen. Some alerts require action depending on the playbook, and SOC analysts will decide whether a higher level of intervention is needed. SOC analysts have different levels, and the lower-level analysts report to higher-level SOC analysts. In certain situations, the analyst must escalate an alert if it doesn't seem normal, leading to action by application owners or administrators. During the crisis exercise, you can test whether the alert and escalation chains are working correctly. The crisis communication strategy can also be tested: what to communicate, to whom, and when.

You should also test whether your backup systems are functioning. Backup procedures must be properly checked, and the business continuity plan (BCP) or business recovery plan (BRP) must be tested. This involves contacting incident response providers, insurance, supervisory authorities, etc., and knowing who to go to under what circumstances. For instance, if you are an SOC analyst, you are part of a 24/7 workforce.

In India, SOC analysts are often not employees of the organization but are provided by companies like IBM (for instance, if you are using IBM's QRadar SIEM tool), or by service providers like Mindtree. These analysts know their supervisory chain of command and their shift times, and they are trained to understand alerts and know whom to inform under specific circumstances. The directory they follow must be accurate, as people may have left the organization, which could impact the escalation chain.

Okay, so now you have to decide the exercise format. So, depending on what is the goal of your exercise—like sometimes we may do this exercise for awareness by the board-level people, or director-level people, or high-level executives—in that case, it will be a different type of exercise than if you want to make all stakeholders aware and all stakeholders debug the procedures and be exact, you know, know how to work together under a crisis and so on. Then the exercise will be much more elaborate, right? So, you have to also see what is your allocated budget, like usually these exercises require some money and also time. Time is money, of course, so if you have to actually make an exercise that is day-long, 2 days, 3 days long, then you are basically losing productivity of employees, so you have to decide on that. So, resource availability for also preparation—so usually, preparation takes many weeks, right, to do an exercise, do an exercise correctly, and the level of experience and, you know, of the people in the organization. Initially, you might do very simple ones and then gradually go for more complex ones.

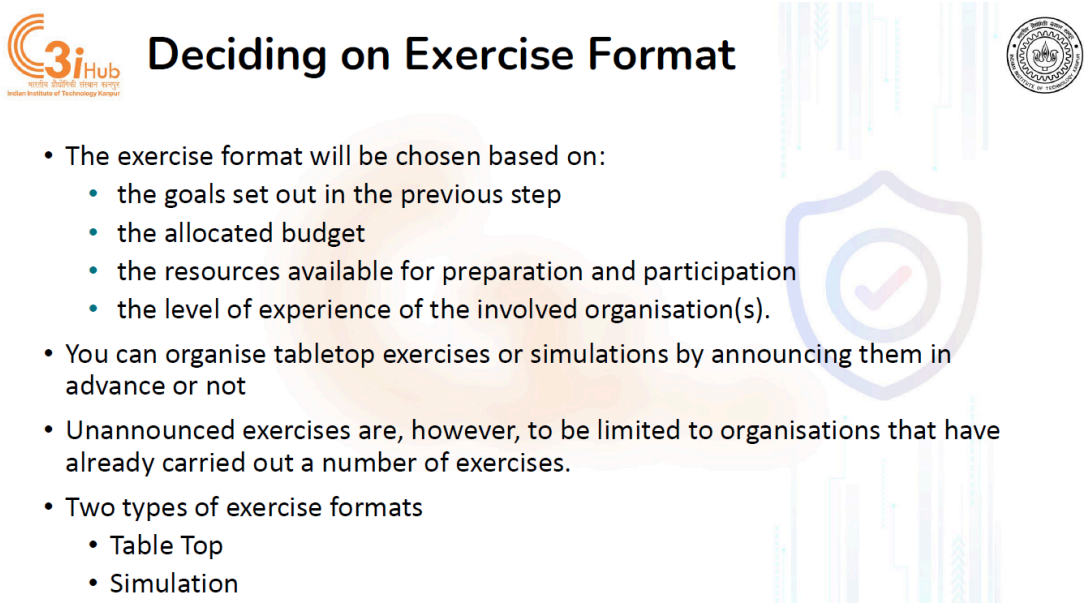

So, the easiest type of exercise that you do is what is called a tabletop exercise. And tabletop exercise is much easier and usually done for top-level management. And then there are simulation exercises, where you do a lot more realistic attack scenarios. Now, sometimes you do, it is the exercise announced, and sometimes you do it unannounced,

right? So, if you do it unannounced, then it might actually create panic in the organization. So, it is recommended by these guys that you do not do unannounced in an organization that has never seen an exercise because that will create a lot of panic. So, the first few exercises, you announce early on.

As I said, there are two types of exercise formats—tabletop and simulation. So, what is a tabletop exercise? So, a tabletop exercise is about decision-making, right? So, you cannot—you are not really going and actually doing the communication or doing the actual, you know, taking out part of the network or shutting down, you know, application and so on. You are just gathered around a table, and then a moderator—there is always a moderator—who would say, who will describe the scenario and ask you to do what we call a thought experiment, right? So, everybody has to come, you know, so everybody will be told to, you know, think of a mindset where a crisis has happened, right, and you have been informed.

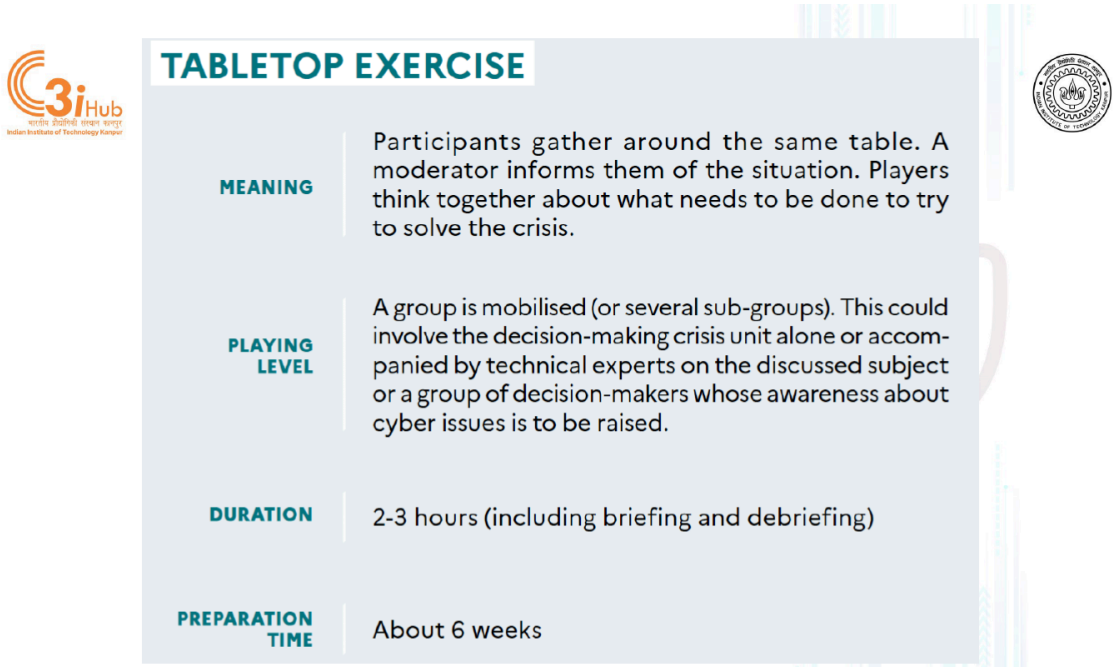

So, at this point, like, you know, it might be that the CEOs and CISOs and CROs, the C-level executives, as well as, you know, the other, you know, major key players of the organization, including communication person, or compliance, legal person, etc., will be around. And then you will be actually saying that, let us say, you have a—you just got a news that, got news that in the morning when people came to work, they found that many of their computers are showing a ransom note and they are unable to log in, etc. And then you can also say after a while, after—so, what needs to be done? Now, the organization

should have a playbook or something that requires to, requires to be, you know, vetted by the organization at various levels, and then has to be adopted, right? So, everybody is given that playbook early on so that they know exactly what they are supposed to do, they are supposed to do. But usually, these incidents are not like, okay, if ransomware happens, this is what you have to do, and if a DDoS attack happens, this is what you have to do, if there is a data leak in the dark web, this is what you have to do. You do not necessarily have all possible scenarios in the playbook, right?

Many times you have to improvise based on your understanding of what needs to be done. And the involvement of the top-level management is important because you have to sometimes make a decision to cut your losses. So, you might actually have to shut down production in order to actually arrest the spread of the threat. But that decision cannot be made by CISO or somebody at the lower level; it has to be coming from the topmost level of management. So, during this thinking, this thought experiment, the CEO will be asked, like, maybe somebody will say, CISO will say that, okay, here is the problem—if we do not shut down production right now, or if we do not shut down our e-commerce website for, like, a couple of hours, we cannot handle this.

Then the CEO has to make a decision right there and then. What needs to be done. So, everybody will have some decision point where they have to make those decisions, and this has to be exercised, right? So, if you do not have the experience of doing this—so, I have had an experience of doing this tabletop exercise, depending on the moderator. So, in our case, the moderator was from Israel, so he knew exactly how to create the tension, right?

So, you have to create the tension, and you have to create the, you know, sense of urgency that this kind of thing requires. And this may actually also involve, like, in case of, let us say, in case of this happening in a power system, large power system company like NTPC and so on. You might have to bring in the regulator also in the tabletop exercise, right? So, the duration of this kind of exercise is 2 to 3 hours, including briefing and debriefing. And in the debriefing, the moderator would say what went wrong, what did not work out, or what was done well, etc. And the preparation time is about 6 weeks. So, it does not happen overnight, like, you know, tomorrow we will do some exercise, and you do that, but you can—you need some preparation. So, the benefits, as you can imagine, it is ideal for an organization that has never done a crisis simulation, and it helps people.

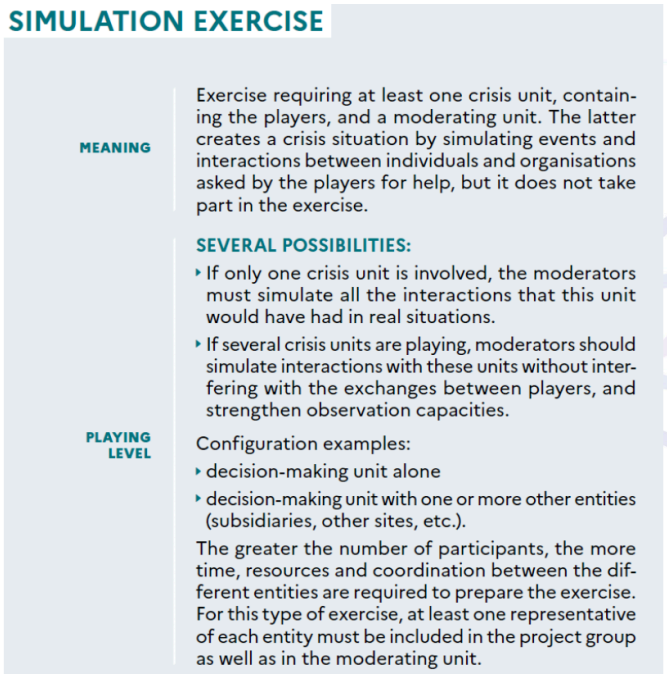

To have a little time to devote to the exercise, so, like CEOs and, you know, C-level executives or board members, they would rather do this than actually do a two-day-long simulation crisis exercise. And it gives you ideas, right? So, you know, once you get into that urgency mode, your mind works towards, like, you know, this can be done or that can be done, so these ideas can be, then, during the debriefing, actually, you know, consolidated and decided whether to accept or not. Now, an actual crisis scheme, like if a real crisis happens, these are all decision-making kinds of simulations, but actually, you may sit and make a decision to shut down this or inform them or do this and that, right? You know, stakeholders and the right set of employees do not do the right thing, then all your decisions are of no use. So, in order to actually do—if in order to check whether your crisis schemes actually would work in the right way, then you have to actually do this in a more elaborate scheme, and that is what the simulation exercise is all about.

Simulation Exercise

DURATION	Between half a day and two days (including briefing and debriefing).
PREPARATION TIME	Around two to six months.
BENEFITS	Deep immersion, increased awareness, makes it possible to test interactions between several crisis units, and shows the strengths and areas for improvement throughout the cyber crisis management scheme.

So, a simulation exercise is actually very long. It is usually between half a day to two days; it may also take longer. So, you have to, in your crisis management plan, have various crisis units, right? So, crisis units—you know, what we were saying earlier, we did not go into the kind of details that you require to actually build a simulation exercise. Your organization should have certain crisis units—so, certain crisis units related to earthquakes, fire, flood, etc., but you can also have crisis units related to cyberattacks, you can have crisis units related to communication, etc. So, you have to actually have one or more crisis units playing, and then there is a moderating unit.

Benefits of Table Top Exercise

- Ideal for an organisation that is not used to running cyber crisis management exercises or that wishes to raise awareness on this topic.
- This type of exercise helps people who have little time to devote to an exercise to get familiar with the topic (management, executive committee, etc.).
- It allows ideas to be drawn up to establish or improve a cyber crisis management system.
- However, it is not possible to test your organisation's crisis schemes with this exercise.

The moderating unit is the set of people who are actually designing the exercise, as well as, during the exercise, they are creating the—narrating the story that is unfolding, as

well as, they are going to create the sense of urgency that is required. Now, if there are—all the crisis units are there, then, obviously, it will be a very involved and realistic one, but in case not all crisis units are at hand, then you can actually simulate the other crisis units' actions by the moderating unit. So, the moderating unit can actually say, I will do the—I mean assume that the communication is done, or assume that, you know, certain network actions have been taken, etc. So, you can have various types of configurations, like decision-making units alone, or decision-making units in, and various other entities, including subsidiaries and other units. So, more participants means a more realistic simulation exercise.

So, it takes about half a day to 2 days, including 2 to 6 months of, you know, preparation time. And the benefits are that it gives you a deep immersive experience, including, you know, increased awareness, including the—you know, much—you know, not only the top-level decision-makers, but it can involve very—you know, real employees who are actually involved in doing things, and also others who might be watching this exercise may actually get some more awareness. You can test the interaction between various crisis units and show where it needs improvement in terms of communication, because multiple crisis units have to work together. Now, what theme are you going to take? Like, are you going to make a theme on a ransomware attack, are you going to make a theme of data leakage, or are you going to make a theme of, you know, an industrial accident induced by a cyberattack, or are you going to make it something like a DDoS attack or malware attack, etc.? So, you have to decide what is the theme that you are going to do.

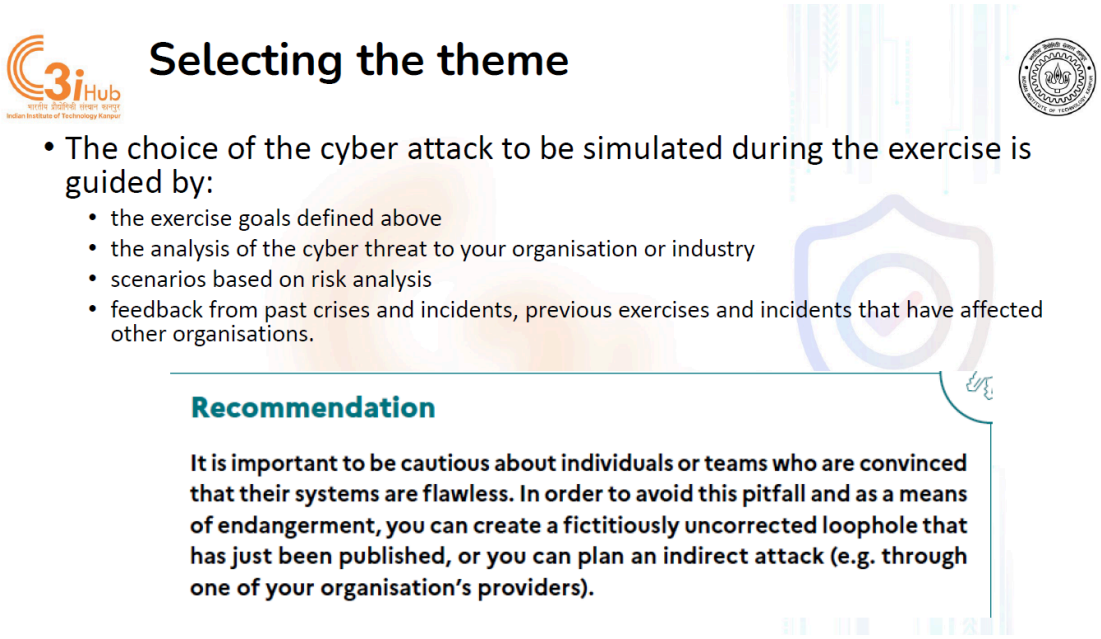

And to decide on the theme, there are several things that people would have to consider. One thing to consider, of course, is what is the threat landscape for that particular sector. If that sector is getting affected by ransomware a lot, then maybe ransomware is one of the exercises that you must do. If it is that, you know, some other, like data stealing or data exfiltration, is one of the major threats that people are facing, then you have to act accordingly. And also, you have to say what are the past crises that have happened, and so on. Now, what this recommendation says in this French document is that some people are often convinced that the systems are flawless, right? So, in order to avoid this pitfall, you may actually do a little bit of a game, right?

So, actually, you can use an uncorrected, unpatched vulnerability and do something to scare people. Or you can also, you know, use your system administrators to actually make it look like a lot of your computers have been ransomed, right? So, a ransomware note will show up on the front page, and so on. You can do certain things like that. But this is a bit of a risky thing. So, the events that you normally assume or create your theme around are, of course, website defacement, which is a simple one.

It is actually more reputational damage than actual harm. But denial of service could actually shut down your business for a while. Data exfiltration—both personal data, like, you know, PII (personally identifiable information), as well as critical data that is critical to your business—may not be personally identifiable. So, you are not under DPDP or digital privacy acts, but you might actually lose business-specific important information, such as IP (business secrets), and so on. Data encryption and destruction, like ransomware attacks, and destruction of services—so, they might actually get into your system, shut down services, and remove the code, and this kind of attack.

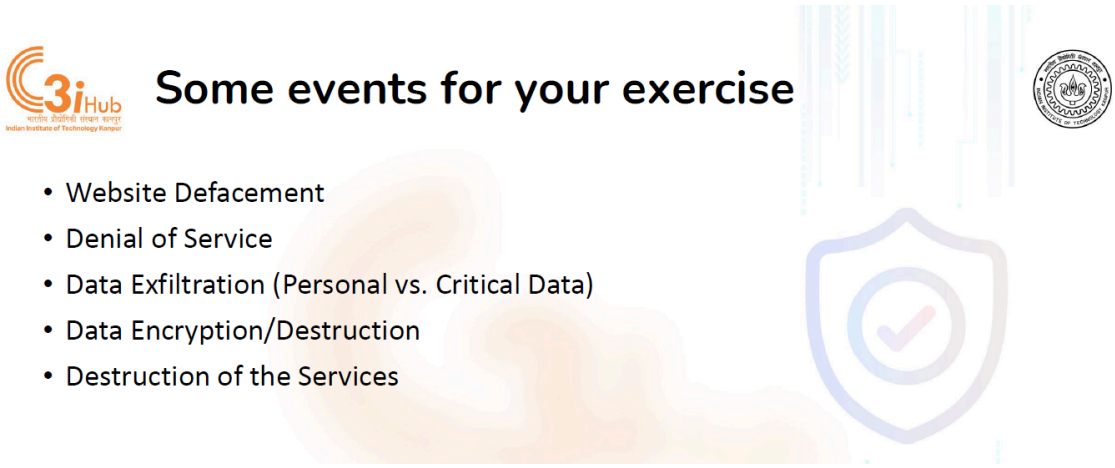

So, these are just a few examples; you can think of other possible events to exercise. The potential impacts considered in the exercise, of course, are: first, reputational elements. Reputational damage costs money because you may lose future customers. So, you have

to take that into account. And also, you have to think about the effect on business continuity, which means that you will lose money by not being able to sell the products or not being able to, you know, produce the products that you are normally producing in your factory, etc., or you are unable to produce power, and hence, you are losing money, etc. You can have reputational impact because personal data is lost, and therefore, you are attracting the wrath of GDPR and DPDP, this kind of regulation. You might have other commercial impacts, including the trust of your customers.

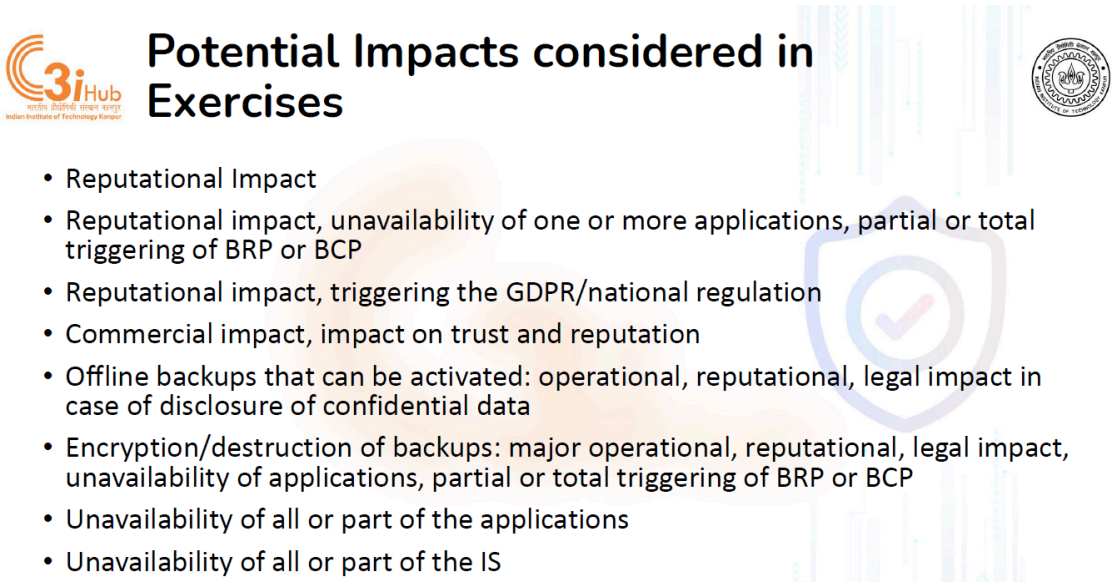

You also have to create scenarios in which backups have to be tested, like whether your backup system worked and whether it has the latest backup. So, banks, for example—based on RBI regulation, all banks have, usually now, three data centers. So, they have a primary data center, or what is called a DC (data center), and then they have a DR, which is their recovery in case the primary data center gets attacked or suffers from any other damage, like fire. They should be able to do all the activities from the secondary data center or DR. And then, nowadays, they have a disaster recovery, which means if both of them get affected, then the third one has to be working. So, in doing so, what happens is that, in a bank, there are multiple applications, usually in the order of 30-40 different, very sensitive, real-time applications, like core banking applications, for example.

So, these applications have to be replicated in all the data centers, and they must always be in the running state, right? So that, in case one fails, there would be a failover recovery in the other data center, right? So, RBI mandates that all banks have to test, every quarter, a real recover failover, right? So, they actually intentionally have to move all their

applications from one data center to the other. And customers should not know—so if you are connecting to, like, SBI net banking or SBI, you know, or whatever, it should be transparent to you that right now, SBI is moving all their applications. So, all the requests from the customer now move to the other data center. So, this kind of thing often happens. So, yesterday, I was talking to a bank, and they had some applications not running in their other data center.

So, during this process, they found that one of their applications—customers could not do this because they had—customers could not be served because they did not have that application running on the other one. So, this kind of test has to be done often. So, in the case of a simulation crisis, you can also create a scenario in which your primary data center has been affected by the attacker, and the secondary one has to be initiated. All the customers have to be moved to the second one, or maybe the third one. And so, every quarter, the banks are not only supposed to test, but they also have to give the report of that test to the RBI every quarter. So, this is the process. The same thing with the exchanges, right? So, for example, securities exchanges like Bombay Stock Exchange and National Stock Exchange—they actually have to do the same exercise, and these secondary and tertiary data centers are in different geographical locations. So that, if there is an earthquake in Bombay, the other one may be in, like, Chennai, right? The other data center may be in Chennai.

So, the chances, the probability that both will have the same natural disaster is low. But in terms of cyberattacks, you cannot make that prediction—they may both be affected, right? So, you have to create that separation, segregation between the data centers, so that, you know, something that is affecting one does not cascade into the other. Okay, so that is all I wanted to say about cyber crisis management. Now I am going to go into the next topic, which is cyber resilience review.

So, cyber resilience. So, we have been discussing this a few times here: cyber resilience is basically your ability to fight through a crisis, right? So, when you have an attack, if your organization completely collapses and does not recover for a long time because it does not have backup, it does not have the right ability to actually rebuild the systems that have been affected, and so on, then we say that the organization has no resilience, right? So, this is the idea of resilience.A city is considered resilient if, when a natural disaster happens, like a hurricane or a storm, if the entire infrastructure—the power infrastructure, the water supply infrastructure, the gas supply infrastructure—if everything collapses and it takes days for recovery, then we say that the city does not have resilience, right?

CS 668:
Cyber Resilience Plan for an Organization

Through a Review Methodology

So, resilience is a term that is used for any kind of system that needs to recover. And not only recover—it should be able to withstand the disaster in a way that it is graceful. It does not completely collapse. It might have reduced quality of service, it might not provide the agreed-upon level of service, but it should be able to at least provide some minimal service. And as time progresses, it should be able to recover without much problem or too much additional, unplanned intervention.

And that time it takes to recover is kind of a measure of what we call resilience. So, it is kind of an elasticity measure, right? So, if you have something elastic, if you deform it, it will immediately come back as soon as you remove the force that deforms it. But if it is not elastic, then it will actually deform, and it will very slowly regain its original shape, right? So, then we say that is less elastic. Same thing with resilience—if you deform the situation, and it comes back to normal very quickly, then we say it is more resilient than the other system where it is not like that. So, cyber resilience is very important. Oftentimes, people, earlier—10 years, 20 years ago—used to think that, okay, we have to stop all cyberattacks, and that is the only thing we have to do, right? So, we have to put firewalls, we have to use antivirus, we have to use whatever means we have to stop the attacker from entering our system.

Unfortunately, that scenario has changed, and we have found that attackers are usually clever, and they often, more often than not, defeat or evade our protection or detection systems. So, we have to assume that attacks will happen, no matter what level of security I provide. That does not mean we should not provide the best security possible, using defense-in-depth and all that. Irrespective of all that, attacks will happen. Even if you have the best perimeter defense, the problem is that there could be insider attacks. There could be social engineering to bring an attack into a system that has been completely secured in the perimeter, right? Also, the perimeter is not well-defined anymore because of work-from-home, multi-site organizations, and so on. So, the process that we have to remember is that we have to assume attacks will happen. So, the question now is that,

when attacks happen, you reduce the probability of the attack happening, but you cannot reduce the probability to zero. Therefore, we have to assume that attacks will happen in the worst case, and we have to be prepared for that. And that preparedness is basically preparedness for resilience.

Resilience is kind of like the property that says that when I am attacked, and my services are compromised, my business continuity is compromised, I do not necessarily completely collapse. I gracefully degrade, but as soon as this problem is withdrawn by the attacker, or I am able to stop the attacker, I should be able to get back to a normal service scenario. Now, what the Department of Homeland Security in the US—they are very concerned—they have this organization called CISA. CISA is the Critical Infrastructure Security Agency within the Department of Homeland Security. The Department of Homeland Security is very concerned because, eventually, all the critical infrastructure entities are no longer government-controlled, right?

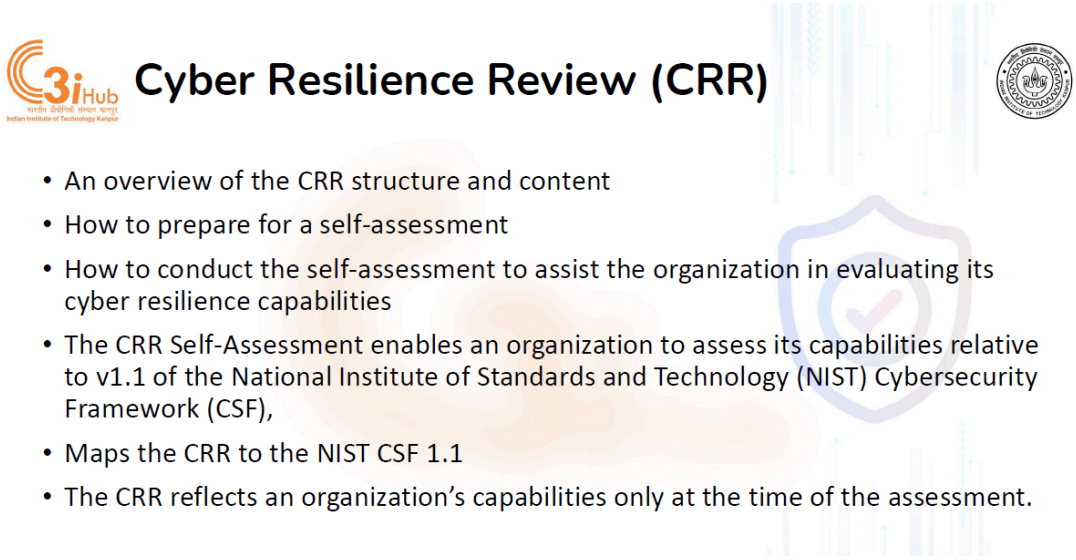

So, even in India, right? You have private companies like Tata Power, Adani Power, and Reliance Power providing power. You have organizations that are providing water treatment and sewage treatment in a city; they might be private, or even if they are not private, they are managed by municipal corporations, which are not directly under the purview of the central government, for example. So now, the question is, how do you make sure that all these organizations are resilient? Because, eventually, if they do not want to be resilient, they will not be resilient, right? If their board says, 'I do not want to spend money on becoming cyber resilient because I do not believe cyberattacks will happen,' then you cannot really do much unless you bring it under regulation. In that case, there should be a regulator, and the regulator will have the ability to ask for reports, cyber audit reports, resilience review reports, and so on. The regulator will also provide

the regulatory instructions as to what kind of things you have to do in cybersecurity—what you need to do. So, right now, as we speak—like this week and the previous week—

the power system in India, power system cybersecurity regulations are being finalized, eventually, right? So, today, we are in 2024. So, in 2020, this regulation was started to be written, and it took, like, 4 years to eventually finalize, right? Now, these regulations will tell you what kind of cybersecurity controls you have to have—whether you have to be compliant with certain standards such as ISO 27001, or do you have to be compliant with a standard like, you know, ISO or IEC 62443, or you have to provide a cyber audit report every 6 months or every 3 months.

They might also tell you that you have to do a risk assessment every 6 months. They might also tell you that, based on risk assessment results, you have to do certain things. They might also tell you some specific controls, such as: you have to have firewalls, you have to have segregation of your network, you have to have two-factor authentication for all employees, you may have to do this kind of network segregation between your plants, where the actual power generation is happening or transmission is happening, and your IT system. So, this is all part of the regulation, right? So, similarly, what these guys have done—they said, 'Okay, the Department of Homeland Security said, "Okay,

let us look at what NIST is telling us."' NIST CSF, or Cybersecurity Framework, and actually, NIST just last week published NIST CSF 2.0, but this thing is not yet according to 2.0. So, we will talk about NIST CSF 1. NIST CSF 1 is a cybersecurity framework that was developed by NIST along with many different stakeholders, including stakeholders in the Department of Energy, oil and gas, and manufacturing, etc. They said, 'What are the minimum things that one needs to do for the cybersecurity of critical infrastructure organizations, right?'

So, that is NIST CSF. NIST CSF tells you that you have to have 5 functions, right? So, they call it Identify, Protect, Detect, Respond, and Recover. So, these are 5 functions. In NIST CSF 2.0, they added a new function called 'Govern,' which was earlier part of all of these.

So, there are 6 functions now in this CSF. So, every organization must check, 'Am I fulfilling this function? Am I fulfilling this function?' and so on. So, all the functions must be there, implemented. The right stakeholders must be identified, and the right technology to implement those functions must be in place. The second thing that NIST CSF says is that you have to have certain tiers of implementation, right? So, not everybody needs to be in the topmost tier of implementation of cybersecurity because, let’s say you are a small, you know, a small-town water treatment plant, right? You are

only, you know, you are only supplying water to a thousand people, right? I am talking about a US scenario.

So, then the question is, how much cybersecurity is required? Now, it has been the case that, in the US in the last few years, there have been numerous attacks on water treatment plants, you know, even in small towns, right? Whoever—it is opportunistic, right? So, it was not necessarily very planned. So, the question is, how much resilience should that have? Now, in a town of 1,000, if water suddenly stops in the middle of the summer, it is going to be a problem, irrespective of how small the town is, right? So, it is actually something—but the tier of implementation may not be the highest tier of implementation for this kind of organization. So, now the question is, how does the Department of Homeland Security know whether an organization is resilient, right?

So, 'resilient'—and remember, resilience is a relative term—because if your business is to actually, you know, supply water, purify and supply water from a lake near you to a town of 1,000 people, then your business continuity is about just keeping the supply on, right? So, it is a very simple organization. Whereas if you have an organization that is supplying water, you know, for the whole of New York City, for example, or Manhattan, for example, then it would be a lot more complex. It might have other things, other than just the continuity of supplying water. So, based on the business—what business you are in and what it means to be serving your customers, what it means to be providing the quality of service and upholding the service level agreement—your resilience will be defined accordingly. So, resilience is not necessarily fixed for every organization, like, you know, what level of resilience is required.

Now, in order for a regulator to know that you are resilient, they have to measure your resilience, right? And to measure resilience, what the DHS, the Department of Homeland Security, has done is that they went to some document called the Resilience Measurement Model, created by the Software Engineering Institute at Carnegie Mellon University, and it is called RMM. And then, using that, they have created a nice questionnaire. And this questionnaire has a very good structure, which we will see. So, you can actually do what is called a self-assessment, or CRR—CRR is the Cyber Resilience Review—so you actually can review your cyber resilience using this self-assessment questionnaire. And then, you can interpret the report, whatever you get, and then you can do follow-up to actually fix the problems that you identify with the resilience problem. So, this questionnaire has 299 questions. So, you have to go into an organization.

3 Phases of CRR Self Assessment

- CRR Self-Assessment Form Completion
- Report Interpretation
- Follow-Up.

What is CRR

- The CRR is a lightweight assessment method that was created by the Department of Homeland Security (DHS) for the purpose of evaluating the cybersecurity and service continuity practices of critical infrastructure owners and operators
- The CRR, consisting of 299 questions, is typically delivered in a six-hour workshop led by facilitators from DHS.
- The facilitators elicit answers from the critical infrastructure organization's personnel in cybersecurity, operations, physical security, and business continuity.
- The CRR Self-Assessment Package allows organizations to apply the same method without the participation of external facilitators.
- It contains the same questions, scoring mechanisms, and options for improvement as the externally facilitated CRR.

CRR Practice

- The CRR is an interview-based assessment of an organization's cybersecurity management program.
- It seeks to understand the **cybersecurity management of services**, and their associated assets, that are critical for an organization's mission success.
- The CRR focuses on protection and **sustainment practices** within key areas that typically contribute to the overall cyber resilience of an organization.
- The CRR measures essential cybersecurity capabilities and behaviors to provide meaningful indicators of an organization's operational resilience during normal operations and during times of operational stress.
- The CRR is derived from the CERT Resilience Management Model (CERT®-RMM), which was developed by the CERT at SEI@CMU.
- The CERT-RMM is a capability-focused maturity model for process improvement, and it reflects best practices from industry and government for managing operational resilience across the disciplines of security management, business continuity management, and information technology operations management.

So, DHS can actually send—so, if you say to DHS, 'I do not have the right people to do this self-assessment,' they will send people to your organization. And they will say, 'Okay, give me 6 hours, and these are the people who should be assembled with me, like people who are involved in IT, people who are involved in security, people who are involved in risk, and so on.' And then they should answer my questions. And once I get the answers, I will fill in the questionnaire, and then the questionnaire will automatically score you, right? So, this is an instrumented file, which can actually—once you answer all the questions, it will actually give you, show you what the resilience level is, where it needs to improve, and all that stuff. So, that is the beauty of this document. So, I will share the document with you; you can play with it a little bit.

So, here is the document. So here, you know, you can answer questions like—so, let's quickly look at some questions. So, it is divided—each—it is divided into 10 different domains, right? So, every organization that needs to be resilient has to have different domains in which it should be resilient. So, here is your asset management. So, do you know what assets, what cyber assets, are actually connected to your network, right? How many servers are there? What are those servers for? What services are they involved in? What kind of operating system, firmware, applications do they run? All that information should be documented and updated as soon as that server has been updated, or a new application has been put into that server, and so on.

So, asset management. Within asset management, they say, 'Okay, here are certain goals that need to be fulfilled by the organization.' For example, goal one is: services are identified and prioritized. See, eventually, why do you have assets or why do you have an IT system? Because you are providing some service. That service might be a customer-facing service, like, for example, if you are Amazon, you are actually running server farms or cloud, etc., to run your front end of your Amazon shop. But you also have services such as goods management, you may have services for employee payroll, you might have some services for HR management. So, you have many services in the organization; some are customer-facing, some are for the actual running of the organization. So, in order to identify assets, you have to first identify what the services are, and then you have to identify what assets—like what servers, what network equipment, what kind of database server, etc.—are involved in providing that service. So, if you can identify all the services, and then you can identify all assets associated with that service, then you can actually figure out whether you have anything redundant, whether there are things that are not even required in your organization.

So, the first question they will ask you is: are the services identified, right? Now, you can answer yes, or you can say no, or you can say incomplete—that is, 'I have started identifying all the services, but I do not have every service identified.' In a very large

organization like Amazon, this might be quite a challenge to actually identify all the services. So, are the services prioritized based on the analysis of potential impact if the services are disrupted? So, how can somebody disrupt a service, right? By actually, I mean, from the cyber angle. From the cyber angle, you can disrupt a service by actually disrupting one of the assets involved, right?

1 Asset Management

The purpose of Asset Management is to identify, document, and manage assets during their life cycle to ensure sustained productivity to support critical services.

Goal 1 - Services are identified and prioritized.		**Yes**	**Incomplete**	**No**	
1. Are services identified? [SC:SG2.SP1]	G	☐	☐	☐	C
2. Are services prioritized based on analysis of the potential impact if the services are disrupted? [SC:SG2.SP1]	G	☐	☐	☐	C
3. Is the organization's mission, vision, values and purpose, including the organization's place in critical infrastructure, identified, and communicated? [EF:SG1.SP1]	G	☐	☐	☐	C
4. Are the organization's mission, objectives, and activities prioritized? [EF:SG1.SP3]	G	☐	☐	☐	C

Goal 2 - Assets are inventoried, and the authority and responsibility for these assets is established.		**Yes**	**Incomplete**	**No**	
1. Are the assets that directly support the critical service inventoried (technology includes hardware, software, and external information systems)? [ADM:SG1.SP1]	G				
People		☐	☐	☐	C
Information		☐	☐	☐	C
Technology		☐	☐	☐	C
Facilities		☐	☐	☐	C
2. Do asset descriptions include protection and sustainment requirements? [ADM:SG1.SP2]	G				
People		☐	☐	☐	C
Information		☐	☐	☐	C
Technology		☐	☐	☐	C
Facilities		☐	☐	☐	C
3. Are both owners and custodians of assets documented in asset descriptions? [ADM:SG1.SP3]	G				
People		☐	☐	☐	C
Information		☐	☐	☐	C
Technology		☐	☐	☐	C
Facilities		☐	☐	☐	C
4. Are the physical locations of assets (both within and outside the organization) documented in the asset inventory? [ADM:SG1.SP3]	G				
People		☐	☐	☐	C

So, either the application, or the network, or the database server, etc. So, the question that then we will ask is that, once you have said, 'Here are the services 1, 2, 3, 4, 5, 6; these are my services, these are the assets associated with the services,' now I have to say, 'Okay, if this service gets interrupted, how much impact will it have on my bottom line, on my business continuity?' If you find that some services I can do without and some

services I have to have running all the time, then the one that is supposed to run all the time has to be prioritized over the one that I can do without, right? So, that is the prioritization of the services. Because without the prioritization, I would not know where to put more security, you know, to protect. Is the organization's mission, vision, values, and purpose—including the organization's place in critical infrastructure—identified and communicated?

Now, this is very important. Now, one might think, 'What has this got to do with cybersecurity?' Interestingly, an organization—especially a critical infrastructure organization—so, now remember, this is designed for critical infrastructure organizations. So, for a critical infrastructure organization, one has to figure out whether my organization and the things I do, like my mission, my purpose, etc., are aligned with the national critical infrastructure. So, if I am a power generation company, like I am like UPPCL, I would need to know what my impact will be if I go down, right, to the entire power system. Because remember the 2012 blackout that was created—it actually started from UP, right? So, basically, one transmission line went down, and then it started cascading. Then it started going to all the—not like—600 crore people were powerless, right? So, it can be one power organization that could actually start a cascading failure. So, you have to know what my impact is if I go down on the overall scheme of things.

And then, are the organization's mission objectives and activities prioritized, right? So, you might have multiple different missions and multiple different—like a power company—but which one is going to be more important? Or, you might have multiple different missions, including one in power, one in gas—you, the same company—then you have to prioritize which one will affect more. So, these questions are to be answered in this way, and I will go through this. You will basically be answering—I mean, not you, I mean whoever has to do this—299 questions, and then you have to basically figure out—it will automatically—see here, they have programs, the JavaScript programs. So, you can actually print the report, you can print the assessment form, and so on. So, this is all automated. Last time, last time I taught this class, I actually asked different groups to actually go to CC and ask these questions. I do not think that went very well. So, but this time, we will not go into that. But you should know why these things are important and what the questions are and why those questions were put into this assessment. So, we will take it from here next class.