How a well-thought-out incident response can take the advantage back from attackers

Received (in revised form): 21st January, 2022

James Christiansen
Vice President, Netskope, USA

James Christiansen is Netskope's Vice President and Chief Security Officer, Cloud Strategy. He is focused on Netskope's global strategy to drive thought leadership in cloud security transformation. James brings extensive expertise as a global leader in information security. Prior to joining Netskope, he was Vice President CISO at Teradata, where he led the global security, physical and information security teams. Previously, James was Vice President of Information Risk Management at Optiv, Chief Information Risk Officer for Evantix, and CISO at Experian Americas, General Motors and Visa International. In each of these organisations, one of the key responsibilities was global incident response. As a sought-after expert speaker on security, James has been featured at numerous prestigious events, including the Business Roundtable, Research Board, American Bar Association, American Banker, the RSA Conference, BankInfoSecurity, ISSA, ISACA, HIMSS and MIS Training Institute. He has also been featured in the *New York Times* and quoted as an expert in *USA Today*, the *Wall Street Journal*, Reuters, *United States Cybersecurity Magazine*, Bloomberg and Healthcare IT News. James is a patent inventor and has received three innovation awards in cyber security, GRC and cloud computing. He is the author of the 'Internet Survival' series and contributing author of 'CISO Essentials', as well as numerous industry papers. He earned his Master's degree in business administration with a focus on international management and his Bachelor's degree in business management from Westminster College.

Cloud Security Transformation, Netskope, USA
E-mail: james.christiansen@netskope.com

Abstract The formation of an incident response (IR) team and IR testing are the most significant actions an organisation can take to reduce the cost of a security breach, but organisations are often challenged to build the right IR team with the right outputs and outcomes decided — especially when IR itself needs an aggressive rethinking in an era of thousands of software-as-a-service (SaaS) applications in use by businesses. This paper will explore how to build the right IR team processes, key roles, tabletop exercises, protocols, executive management and other important considerations. It will highlight both hard and soft skills needed for successful IR, especially in the less-discussed but hugely important latter category, the need for expectations setting with leadership, and the right ways to convey vague and incomplete information in the early stages of IR and breach analysis by using actual cases from past experience.

KEYWORDS: incident response, IR, SaaS, cloud security, cost of a breach, cyber security

INTRODUCTION

The average total cost of a data breach increased by nearly 10 per cent from 2020 to 2021 — the largest single-year cost increase in the last seven years.[1] The formation of an incident response (IR) team and IR testing are the most significant actions an organisation can take to reduce the cost of a security breach, but organisations are often challenged to build the right IR team with the right outputs and outcomes. This paper will discuss some of the most important considerations for IR in the hybrid cloud era and provide practical advice on how to achieve the right outcomes.

BUILDING AN EFFECTIVE IR TEAM

This paper is focused on how to plan and train for an incident to take away its natural advantages. The first step is to establish the roles of the different team members and how they will perform during an incident. For starters, designate a team leader, who is most often the CISO or their delegate. IR will also typically include members from the legal department — senior staff who will monitor the situation to ensure that legal procedures are followed and work with external counsel as needed. The team will need a business leader to understand and advise on the business impact of the event. They typically will work with corporate communications to keep the business units informed about what is going on and how it is affecting them. The team will also need a customer service representative — as customers are calling in to say systems are down or enquire about the event once it becomes public, the customer service team knows how to respond. Technology leaders are the team members who are actually trying to recover systems. A security leader will be working on how to contain the incident to minimise the impact, doing attack analysis, forensic analysis, eradication and supporting any e-discovery as needed.

Once the team is assembled, the first order of training should involve going through rules with the incident response team concerning who does (and does not) talk about the incident to anyone outside the company. Normally the corporate communications team, which typically includes public relations, is designated as being the point of contact for all external communication, and everybody else needs to know that they cannot make any kind of comment without explicit approval from corporate communications. Even if they are asked by a neighbour or friend, everyone must be trained to defer comments to the designated contact person. Best practice includes a special non-disclosure agreement (NDA) with the team members to reinforce the need for confidentiality.

TABLETOP EXERCISES

Once an incident is in motion, things can get chaotic. This is where training and tabletop exercises will pay off. The IR team will use tabletop exercises to simulate actual attacks and work through the process and decisions that need to be made along the way. Best practices include tabletop exercises for the technical staff and executive leadership. Oftentimes, lack of a clear game plan that anticipates the cycle of a response effort is what gives hackers added advantages that compound the costs and damages of an attack. Whether through lack of resources or inadequate preparation, there are many ways that organisations can tip the odds of a breach in favour of cybercriminals. These are some common problem areas that can help IR team leaders stay one step ahead during a crisis.

Best practice is to create a set of scenarios that represent real risks to your organisation. Establish a tabletop exercise that includes an overview of the scenario, then the steps to proceed. The incident response leader will provide bits of information to the incident response team. The team then discusses the reactions to the information provided and the steps each member will take to gather additional information for analysis or begin containment of the situation. The team leader should record the outcome of each step to later be used in a post-mortem review of the exercise. The scenario continues to unfold as the team leader provides additional information.

It is common in an actual incident for the information to become increasingly bad as the scenario goes forward and analysis points to a more significant breach than first imagined. A key point at each step is to discuss what communications should be made to the management team, customer service team and staff. There is a delicate balancing line between the need to provide status and impact information and reporting too soon before the impact analysis is complete. It can often take a great deal of time to determine the total impact and providing revised results can be costly in public opinion.

Once the tabletop exercise is complete, do a review of the key issues identified in the scenario:

- Was the communication smooth between the IR team, management and key stakeholders?;
- Did the team have the necessary skills to resolve the incident?;
- Were additional outside resources available in a timely way to help resolve the problem?;
- Was the scenario realistic, complete and reflective of an actual threat to the organisation?;
- Were any additional team members identified who should be included in future tabletop exercises?

Best practice is to perform a tabletop exercise at least quarterly with the technical team and at least biannually with the management team. Tabletop exercises include phishing privilege account takeover, ransomware/extortion, advanced persistent threat and widespread virus attack, cloud SaaS application. In addition, consider scenarios that directly relate to your business and the unique infrastructure and product delivery processes.

Now let us look at some scenarios that point out common failures in incident response that give the advantage to the hackers.

NOT KNOWING WHAT'S NORMAL

The hacker will have an advantage if the company does not understand what is normal in its infrastructure. If one does not understand what is normal, one cannot understand what is not normal — this means being able to distinguish if something is an anomaly. If certain traffic looks different, maybe there is something wrong with it. This kind of ongoing awareness can help detect an attack early on and potentially save the organisation from a major security event.

For example, maybe the business is getting reports of inoperative systems with a message about paying a ransom. Calls start coming in from users who are unable to access the system. An attacker has encrypted a server and they are giving 24 hours to pay US$1m. And if the company does not pay by the deadline, the ransom increases to US$2m. That price will keep escalating the longer it takes to make the payment. This is an increasingly common situation for businesses of all sizes. Over the last year, ransomware attacks have grown 150 per cent and the average amount paid by victims has tripled.[2] The average total cost of a ransomware breach is now US$4.62m.[3]

To prepare for a potential ransomware attack, IR teams need tabletop planning sessions, a clear understanding of all assets and a solid backup plan. The key is being prepared: understanding where the organisation is from a cultural perspective, a business perspective and an individual perspective. With those details out of the way, the IR team leader can focus on the specifics of an actual attack when it happens. What business processes are out of service and what is the impact? Is the organisation open to paying or is it doing a recovery? To make those decisions quickly, the IR leader needs to have a ransomware plan already in place. There may be an 'oh no' moment: you find out that recovery from backup will take over seven days. What is the impact on the business now? Should we consider paying the ransom?

This is a controversial subject, and for good reason. Security professionals never want to see anyone have to pay a ransom. But it is still a good idea to have a corporate Bitcoin account set up in advance, in case of the worst-case scenario. At the end of the day, a business decision needs to be made, gauging the impact of the outages and potential data exposure against doing system recovery. Hopefully, thanks to the tabletop exercises, the ransomware is contained to a small exposure that can be eradicated and normal business recovery without the need to pay the ransom. Every organisation will need to plan for and decide if, and when, it is time to pay; the situation and its exact protocol cannot be made 'one size fits all'.

Here is another example of not understanding what is normal in the infrastructure. Let us say the company is getting a lot of anomalies in its security reports — so many, in fact, that it is getting flooded with information and the IR team does not know which issues they should investigate. I have seen this sort of thing happen frequently in past breaches; the alerts were there, but nobody could spot the good information within the overwhelming volume of bad information.

The IR team leader needs make sure they are filtering in advance to sift the relevant information. The IR team needs to take swift action using good data, rather than waste time triaging false positives in security reports. The IR leader has to keep the security team operational at all times. If they are not responding to an immediate incident, they can use that time for proactive tasks like threat hunting that help reduce the likelihood of an incident in the first place.

NOT PREPARING IN ADVANCE FOR ATTACK REALITIES

Let us say an organisation has some systems down as the result of an attack and the IR team lead needs to contact the suppliers of some of these technologies. Maybe they need some rules set, or possibly even some additional support. The problem is that with system outages, no one can access the vendor support numbers. This becomes a real issue, because time is an absolutely critical factor.

In the middle of an incident, IR teams need to be able to react fast and make decisions with urgency. To do that, they need to have everything organised to anticipate even the most basic problems. That means having all emergency numbers readily accessible on a local resource, because if, for example, it is stored somewhere on a server that has just been wiped out, that is not going to be any use. Team leaders may need contact info for security vendors to help contain this very incident, or they might need contacts for internal staff escalation or notifications out to the executive team. Backup contacts may even be needed to be used for notifications to incident response team members.

There are a lot of good cloud resources out there for backups like this. Even if local systems are down, it is unlikely the cloud is down, or resources could potentially be accessed from a personal system, if work systems have all been disrupted.

Admin levels under attack

Attacks on admin levels are a very common category of incident. Many organisations set the same admin password for all their systems, and then they share those admin systems because there are maybe only ten techs working across 100 servers. Very often, organisations use a common admin-level password so that whoever is assigned to work on a particular server has quick and easy access to get the job done.

Of course, a compromised admin password offers the highest privilege to an attacker. The pitfall is that if a company uses a common admin password, hackers can quickly move laterally across all the different servers, and it is not uncommon to see that admin passwords have been compromised.

Broadly, let us say the organisation does have good privileged access management in place. In the event of a successful attack in this situation, the attacker has probably compromised an admin password, but the organisation is not going to know which one, even if the passwords are all unique. So the key thing here is being able to quickly change all those admin passwords and communicate the fact they have changed to the support teams. As soon as the passwords change, support teams will not be able to get into those systems. To manage this, organisations need to have a good emergency change control process in place. This process must enable IR teams to make changes quickly, but also communicate out, so that essential people are not locked out of the system and operations are not disrupted.

Playing the guessing game

Time is on the attackers' side. It is very common in the middle of an attack for the IR team to be asking the team leader for a decision when they need to know what actions to take. They may be asking if they should shut down a segment of the network, or whether or not to take a compromised server or specific application offline.

When the IR team leader is being asked to make a decision in the heat of the moment, the result of that decision is typically going to have an impact on the business. If a server is taken offline, the team may lose forensic data that is needed later in order to analyse where the attack came from and determine what (if any) data was exfiltrated. They are going to need the log files to do any of that. In the midst of an active attack — where the business may be haemorrhaging sensitive information — leaving a system online long enough to capture forensic data is always a trade-off. Leaders are often forced into a position to make a guess about what action to take and guessing means they do not have control of the situation.

A related problem is information is coming in from the tech team that is not clear or certain. At the same time, the executive team is asking for answers — answers the IR leader in most cases does not have yet. Remember: any best-guess decisions made in minutes are going to be scrutinised for days afterwards. This is where tabletop exercises will pay dividends both from a planning perspective and in educating other team members on the complexity of the situation and the fact that decisions may need to be made without full information. These are what security leaders like to call 'perfect decisions made from imperfect information'. What was the rationale? What was the impact? IR team leaders have to be prepared to translate the nature of what is happening (including the early-stage uncertainties they are facing) up the chain of command.

This is where planning and experience really come into play. Running legitimate hacks against the infrastructure before the fact can help IR teams understand what might be happening in different attack scenarios. Tabletop training exercises can help IR team leaders learn how to quickly triage which decisions need to be made right away, which ones require further clarification and which ones can be postponed. These exercises can also train and test leaders to communicate clearly, succinctly and unambiguously in the heat of the moment.

Tabletop exercises that test the ability to detect and respond to an attack help IR leaders simulate the pressures of split-second decision making under different circumstances. This is a critical part of an IR plan, because any mistakes can give huge advantages to the attackers. The longer the situation goes on, the more damage they can do and the more data they can exfiltrate.

RUNNING OUT OF ENERGY BEFORE THE BAD GUYS DO

Another common issue in IR has to do with fatigue. When a team is first responding to an active threat inside their system, there is a lot of adrenaline. Everybody is running around, trying to understand where the bad guy is, what they are attacking and how to get it contained. There is typically a lot of emotion and people start at a sprint, but the IT team leader must quickly reframe the situation for what it really is — a marathon — in order to keep people working at a pace that can be maintained for an extended duration. Team leaders have to allow people to rest at about the ten-hour mark, otherwise serious mistakes can be made due to human fatigue. Someone may type the wrong address on the server, or they take the wrong server down, or they enter the wrong passcodes and get locked out.

IR leaders need to have a plan in place ahead of time to make sure operations can continue while team members take turns getting some rest. One common approach is the 16/6 rule: for every 16 hours worked during an incident, IR team members have to get at least six hours off. One thing I always do as part of my scenario planning is locate a nearby hotel to send people over for a few hours of rest in rotating shifts. It is a good option to have at the ready, especially early on in the attack cycle.

Leaders may need to mandate that each member of the team is offline for at least six hours every day. Here is the reason, and it is not always obvious: people on an IR team are not going to want to stop. They are going to want to run that next skip script or do that one next thing on their list, especially if they are making progress or feel 'on a tear'. The IR team leader may have to enforce off-time requirements to keep people rested, sane, sharp and safe for as long as it takes to resolve the problem. These situations can go on for several weeks. Leaders do not want to compound matters with an error made due to exhaustion — or much worse, such as if someone on the team is involved in a car accident on the drive home after too many hours on the job.

NOT HAVING THE NECESSARY SKILL SETS ON THE TEAM

It is most common to think about hacks of on-premises services, but let us say someone has compromised an organisation's cloud-based application. A lot of critical applications are now being ported out to the cloud. While an attack against a SaaS application would be handled by that vendor, a private application running out on Amazon Web Services (AWS), Microsoft Azure or Google Cloud Platform is the organisation's responsibility under a typical shared responsibility model. Many organisations run into problems because no one on the IR team is fluent in cloud. They do not know how to do analysis in the cloud. They do not know how to pull the logs down. They do not know how to run the scripts. All their investigation tools are set up for on-premises investigations. Not having the proper skills and tools within an organisation to handle an attack gives the hackers a significant advantage.

The key thing here is to know the organisation's environment and make sure the team includes the requisite training and skill sets. The shared responsibility model of securing anything placed in the cloud (including applications) is table stakes for any IR team today — including how to handle cloud-based attacks.

I once saw a company go 48 hours trying to actually copy an entire server from the Internet down to local so they could run their analytics. Waiting 48 hours in the middle of an incident is a lifetime. They could not even get any information because they had not put their incident response team tools out in the cloud by the data, and there was significant fallout for the IR and security leaders after all was said and done. This was unfortunate — it was also avoidable.

THE LATEST CRITICAL APPLICATION

The average number of critical SaaS applications being used by organisations is increasing exponentially. Your tabletop exercises should include how to handle an incident in a cloud SaaS application.

Think about it: a corporate security team's response to a breach of a SaaS solution managed by a third party is, necessarily, far different from its response to an attack on the corporate network, data centre or endpoints. That is not just because the security team has less control over the SaaS environment, but also because it has less information that would crucially help to inform the response. Effective attack response requires a series of steps: identification, containment and eradication of the threat; recovery; and learning lessons from the experience. But identifying, containing and eradicating requires deep access and visibility into the software system and its log files.

Now, consider each of the SaaS applications that your company relies on. How much visibility do you truly have into those systems? With that in mind, what would an incident investigation look like if that SaaS application experienced a breach? If you do not know, it is time for that to change. The average company has around 1,200 cloud providers. Some of those are infrastructure-as-a-service (IaaS) or platform-as-a-service (PaaS) providers, which tend to give customers more control over their security environment. But around 90 per cent of the typical company's cloud vendors are SaaS providers. This means the average company's security organisation is responsible for protecting data within more than 1,000 applications, over which they have minimal control and limited visibility.

SaaS applications have become too important to business operations for security teams to let this situation continue. We need to make sure we are involved in the selection, architecture and design of every SaaS solution our organisation uses. We need to be thinking constantly about how to respond to a SaaS security breach. We need to adapt our on-premises incident response plans to fit the SaaS world, modifying our runbooks for different SaaS solutions. And as a profession, we need to insist that our SaaS vendors provide the information we need to effectively manage our company's cyber security risk. No IR training is complete without a detailed assessment of SaaS. For critical and sensitive applications, organisations should also determine if log files can be made available directly to the organisation on an ongoing basis (monitoring) or in the event of a security incident.

MANAGING UP

Another thing that IR team leaders have to deal with in any incident is managing leadership and their expectations under very stressful circumstances for the broader business. In any active incident situation, executive leadership is naturally going to want answers on a continuous basis. Why are the systems down? Are they still attacking us? Has data been exfiltrated? How many records have been compromised? Are they still in our systems? How did this happen? Who do I blame? Who gets fired?

Business leaders need to understand ahead of time that they are likely to receive a continuous flow of bad news early on. In the first days of an incident, team leaders are not going to be giving them any good news. They will be telling them that the attack has penetrated the system. The attackers have exfiltrated data. Systems have been encrypted by ransomware and the attackers are demanding this much money before they send the key. So a good IR team leader will build a relationship with executive leadership well in advance to prepare them for this inevitable part of the response cycle. IR leaders have to educate management and executive teams about these scenarios to help them understand that during the initial hours and days of an incident, there is often not going to be a lot of solid information.

The first reports are probably going to be bad news until the tech and security teams understand the scope of the situation and can actually start getting it turned around with containment and eradication. But that takes time. Sometimes the bad news is actually a good thing, because at least they are getting some clearer answers.

I once managed an international incident where I was on a single call for 16 hours, and I had two different phone lines open throughout. On one line, I had all the tech people who were trying to work the incident. On the other, I had a call with the executive team talking about what we should do. What kind of notification should we be giving? How do we get operations running again? What is going to be the financial impact on the business? I could not let those two call lines merge, because once the executives start asking the tech guys direct questions and demanding answers, then the tech team is no longer working to resolve the incident, they are trying to field questions they do not have answers to yet. Even worse, they are not trained to give executive answers. They will give the executive team what will be perceived as bad answers that will make everybody go crazy. An IR leader has to be the liaison. They are responsible for taking the known information and translating it into business outcomes for management, while keeping the technical team focused and undistracted.

PLANNING AND TRAINING GIVES IR TEAMS THE ADVANTAGE

Threats today evolve so quickly in terms of speed and sophistication, they do not need any extra help. Building a team of cross-functional experts is the first step in effective incident response preparation. But ongoing training and rehearsal exercises are what help IR teams keep a sharp edge for the inevitability of an incident. Once an attack is detected, avoiding these common mistakes can help minimise the time and cost associated with an incident.

It can be very difficult to stay calm in the middle of an incident. But that is maybe the most essential thing that a team leader brings to their team in a crisis situation. It helps everyone else do their jobs better. The best way to maintain composure and control of the situation is knowing that the team has done everything it can in order to be ready for this moment before it happens.

References

1. Ponemon and IBM (August 2021), 'Cost of a Data Breach Report 2021', available at https://www.ibm.com/security/data-breach (accessed 17th February, 2022).
2. Sharton, B. R. (May 2021), 'Ransomware Attacks Are Spiking. Is Your Company Prepared?', Harvard Business Review, available at https://hbr.org/2021/05/ransomware-attacks-are-spiking-is-your-company-prepared (accessed 17th February, 2021).
3. Ponemon and IBM, ref. 1 above.