GIR Global Investigations Review

How to develop and implement an incident response plan to mitigate impact of cyberattacks

 Evan Norris

 Dean M Nickles

CRAVATH

31 July 2025

Business organisations are subjected to attempted cyberattacks on an effectively constant basis. Preparations to withstand those attacks that will inevitably slip through an organisation's network defences can significantly mitigate their impact. One of the most important preparatory steps an organisation can take is to implement an incident response (IR) plan that has been tested for effectiveness.[1] To be effective, an IR plan must be tailored to the organisation and practical, and must reflect institutional buy-in and ownership. An off-the-shelf plan, or one that is not regularly reviewed and updated, is worth little and may even be counterproductive in the event of a crisis.

In this chapter, we discuss the elements of an effective IR plan, including the structural and foundational considerations that organisations should bear in mind when developing or revising their plans.[2] We then provide observations regarding the steps organisations should take to ensure their IR plans remain effective over time. We primarily focus on providing generally applicable advice, but in a few instances flag jurisdiction-specific considerations.

Developing an incident response plan

Scope and applicability

At the outset of IR plan development, an organisation should set out the strategic elements of the plan, including its scope and applicability. While it may seem simple, to build a functional plan, an organisation should decide what it intends to address in its plan, as well as to whom and under what circumstances the plan applies.

For example, an organisation may decide that a particular plan should cover the IR process for a subset of the many functions involved in incident response (such as information security, information technology, communications, legal and compliance).

And, depending on its size and complexity, an organisation may decide that a single plan is appropriate, or it may decide that each business unit requires its own plan. To ensure ownership and the necessary perspective to achieve the intended scope, it is also important to identify who within the organisation is essential to the plan's execution and maintenance. Often, organisations designate their chief information security officer (CISO) and the CISO's team as primarily responsible for the IR plan.

Thinking through and answering these questions can anchor the IR plan within the organisation, align stakeholders on the plan's purpose and guide the approach to the more granular details that compose a plan. Plans crafted without these questions (and answers) in mind can become unwieldy and unfocused, and therefore likely useless during an incident. In addition to serving as a drafting guide, the answers to these scope and applicability questions can also provide important context to readers of the IR plan who were not involved in its development, and should therefore be included at the beginning of a plan.

Definitions

Defining key terms is necessary for development of an IR plan that is well calibrated to an organisation's needs. Among the most important terms to define are those related to different 'levels' of unauthorised, malicious or unknown network activity. An organisation's cybersecurity team will likely be alerted daily to hundreds, thousands or even more potential malicious and unknown activities on its networks and systems. An organisation cannot realistically activate its full IR plan and process for each such alert. It is necessary to differentiate between activity requiring activation of the plan and activity that appropriately can be handled by the information security team and does not require plan activation. Organisations may choose their own terms for these 'levels' of activity, but one option is to adopt the nomenclature of the National Institute of Standards and Technology (NIST), which uses the terms 'events', 'adverse events' and 'incidents'.[3]

In conceptualising these terms, it is important to understand that 'events' will occur much more frequently than 'adverse events', and 'adverse events' will occur much more frequently than 'incidents'. Imagining a classic dart board, 'events' would make up nearly all of the board, while 'adverse events' would make up the outer bullseye and 'incidents' would make up the inner bullseye.

Organisations may choose to modify these definitions to fit their specific needs, including legal or regulatory requirements. A public company, for example, may choose to define 'incident' consistent with the definition of 'cybersecurity incident' set forth by the US Securities and Exchange Commission (SEC).[4] Once those terms are defined, the organisation can evaluate when, and to what extent, the plan and its processes are activated in response.

Other important terms to define include terms with potential legal or regulatory significance, such as 'personal information' and 'breach'. These terms can have significance for triggering notification and disclosure obligations to individuals and regulators, as well as contractual counterparties. Once an organisation has identified

these key terms, the IR plan's usage of them should be consistent with, or at least mindful of, their applicable legal, regulatory and contractual meanings. While using these terms can be challenging during the development of a plan, it can be even more challenging during an incident if an organisation has not prepared itself carefully. Considering these issues in advance can help an organisation protect itself from inadvertently triggering legal, regulatory and contractual obligations and potentially missing deadlines as a result.

Starting with 'personal information', many breach notification laws define 'personal information' (or some analogous term) by reference to an enumerated list of data characteristics considered sensitive. For example, many US state breach laws narrowly define 'personal information' as an individual's first name (or first initial) and last name combined with one or more other data elements, such as a social security or driver's licence number.[5] By contrast, California applies a broader definition, covering 'any information that identifies, relates to, describes, or is capable of being associated with, a particular individual', including a non-exclusive list of identifiers.[6] Outside of the United States, many breach notification laws feature even more expansive definitions of personal information that cover any information relating to natural persons. For instance, the European Union's General Data Protection Regulation (GDPR) broadly defines 'personal data' as 'any information relating to an identified or identifiable natural person'.[7] Because notification obligations are often tied to compromise of 'personal information', organisations should identify the term's definitions in their primary jurisdictions and develop their plans with that definition in mind.

Turning to the definition of 'breach', depending on the jurisdiction, the definitional elements typically include one or more of the use, disclosure, acquisition of or access to data through illegal or unauthorised means. Many US states define a breach as the unauthorised acquisition of personal information,[8] while others only require unauthorised access, and yet others require that the unauthorised access compromise the security, confidentiality or integrity of the personal information.[9] Outside the United States, the breach definition tends to be broader. The GDPR defines a data breach as any 'accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise processed'.[10] As incidents increase in sophistication, the definition of a data breach continues to evolve to include a wide range of activities in addition to acquisition, access, use or disclosure, and organisations should remain mindful of these developments when crafting their IR plans.

Given the nuances of these definitions and their importance, an organisation should be thinking of these terms early in the development process. Because there may not be a simple approach that harmonises definitions across all relevant jurisdictions, an organisation may elect to favour the broadest definitions when developing the IR plan. And because the law surrounding data breaches continues to evolve, organisations

should ensure that they also review and confirm that definitions and key terms in the plan are up to date during their annual review processes, as will be discussed below.

Incident classification

Another foundational element of IR plan development is the creation of a classification system for incidents that enables an organisation to evaluate and respond to incidents appropriately. What constitutes a particular severity level will differ based on the organisation's business. Organisations should consider the qualitative and quantitative factors resulting from an incident that present increased risk to their business, and then assign different severity levels depending on such risk. For example, compromise of intellectual property or trade secrets may present severe risk to a tech company, but less severe risk to a healthcare provider. Organisations can then set up a classification structure to classify incidents as different severity levels (e.g., Severity One, Severity Two, etc.) through the use of those factors. Thinking back to the bullseye analogy, this process represents a further splitting up of the 'inner bullseye' of incidents into severity categories.

Ultimately, incident severity flows through an IR plan and controls the organisation's response. For example, the theft of an employee's work phone and a ransomware attack on a customer database may both result in the compromise of personal information, but an organisation may reasonably decide that the response to the former should be different from the response to the latter.

One nuance to be aware of in classifying incidents is how to approach 'potential' versus 'actual' incidents. Classification decisions are invariably made with limited or imperfect information before an incident's impact has fully been felt. Organisations should think through how, if at all, the plan should treat potential versus actual incidents, and whether a mechanism is required in the IR process for elevating (or lowering) incident severity levels as more information is learned. The theft of an employee's work phone, for example, may initially be classified as a lower severity incident, but that could change if, for example, the thief was able to use the phone to access sensitive company files or the broader company network before the theft was reported.

Legal, regulatory and compliance requirements

In addition to the definition-related legal and regulatory issues discussed above, it is important to identify the key substantive legal, regulatory and compliance requirements related to cybersecurity and privacy applicable to the organisation when developing an IR plan. The plan should summarise such requirements and note where details can be found, for example, in an appendix to the plan or a supplemental, separate document maintained by an organisation's legal or compliance function. For example, an organisation subject to the US Health Insurance Portability and Accountability Act (HIPAA) might include a reference in the plan to the HIPAA Breach Notification Rule,[11] and then include as an appendix a more detailed explanation of the rule's requirements

and how the organisation approaches breach notification in the event of a breach of unsecured protected health information.

Personnel

In addition to the strategic and foundational elements described above, an IR plan should address a few key categories of information: personnel, process and communications.

First, a plan should identify and set forth in detail the roles and responsibilities of the personnel who will be involved in the IR process. The core team involved in the process is commonly referred to as a cybersecurity incident response team, or 'CIRT'. The ideal structure and staffing for a CIRT will be unique to every organisation, depending on the organisation's size, geographic scope, staffing structure and personnel, among other factors. As the plan is being drafted, the organisation should consider the best structure for the CIRT and evaluate and identify the individuals best equipped to fill each role, in consultation with the organisation's information security leadership and the internal stakeholders identified as part of the development process. Once the decision on how to structure and staff the CIRT has been made, the IR plan should identify the CIRT personnel and their roles and responsibilities.

For example, it is common to assign a CIRT lead, who will be the individual responsible for leading the response to an incident, assigning responsibility to team members and reporting the details of the response to senior management. The CIRT lead may be the organisation's CISO, though it need not be. In addition to the CIRT lead, the plan should identify the CIRT members, as well as any special roles in the CIRT, such as a record keeper, those responsible for coordinating with the organisation's IR vendor and those responsible for communicating with other functions within the organisation.

The organisation should also consider which other functions should be involved in the IR process. Leaders in operations, finance, legal, compliance, communications and investor relations may all have roles to play in responding to an incident and may require reference as key personnel. Sometimes, an organisation will have a representative from each of these teams belong to a 'crisis management team', which may be convened for incidents meeting certain severity criteria and thus requiring greater organisational visibility and resources. For each team and individual included in the plan, it is important to clearly identify team leads and decision-makers, describe responsibilities pre-incident, during an incident and post-incident and assign reporting lines. Contact information for all such teams and individuals should also be included in the plan, whether in the body or as an appendix, including communication methods not reliant on organisation systems. If communications systems are down during an incident, having the contact information for all relevant individuals in hard copy within the plan can be vital.

Further, during the development of the IR plan, an organisation should consider the external parties it might need to work with or contact during an incident response and identify those parties in advance. These external parties could include, for example, an incident response vendor, a ransom negotiator, outside counsel, cyber insurance contact

and a public relations firm. An organisation should also consider identifying relevant law enforcement and regulatory contacts, particularly any law enforcement agents, that can help mitigate the immediate impact of an incident. As with internal teams and individuals, contact information for these external parties should be included in the plan.

Incident response process

Second, an IR plan should describe the IR process itself, which can generally be divided into four phases: (1) preparation; (2) detection and analysis; (3) containment, eradication and recovery; and (4) post-incident activity.[12] It is important that the plan describe this process, as well as the steps the CIRT should take and procedures to be followed for each phase of the process covered by the plan, depending on the type and severity of the incident.

From an IR plan perspective, the most important elements of the IR process are often Phases Two and Three. It is vital that Phase Two (detection and analysis) flows logically into Phase Three (containment, eradication and recovery). In other words, the identification and classification of an incident, consistent with the definitions and classification categories decided upon by the organisation, must result in a clear direction for the steps to be taken to contain, eradicate and recover from the incident. Sufficient detail should be provided such that the organisation, including the CIRT, understands how it should handle incidents of differing levels of severity. Organisations should also give thought to their desired process for concluding an IR process during the transition between Phases Three and Four.

The exact level of detail to provide for each phase of the IR process will vary greatly depending on the organisation. The key is for the process described in the plan to be implementable and consistent with the organisation's desired practices, while providing sufficient guidance to be helpful in maintaining uniformity in approach across incidents.

Communications

Third, the IR plan should address internal and external communications during and regarding an incident. The plan should address the proper flow of communications, including the individuals or teams responsible for such communications and when they are to occur. Communications need not necessarily be addressed in a standalone section; this subject can be addressed as appropriate throughout the IR plan or even in an appendix or a separate document, if the flow of communications is addressed and responsibility is assigned. Depending on the organisation's industry and applicable laws and regulations, it may also be helpful to incorporate references to notification obligations as well, such as in the HIPAA example provided above. Organisations also may want to assign specific responsibility for notifying senior management and the board of directors in certain situations, such as particularly severe incidents or (for public companies) incidents that may require a materiality analysis for purposes of SEC disclosure. Incident response is fast-moving and stressful, and establishing a

communications plan can help ensure that information is appropriately communicated across the organisation and to external parties.

Other elements

There are numerous other elements that an organisation may want to include in an IR plan based on its specific circumstances. For example, a public company subject to the SEC's disclosure requirements might want to incorporate the concept of a materiality assessment into its IR process. For example, when classifying the severity of an incident, the IR process could require that incidents of certain severity levels be subject to further assessment for materiality-related considerations and escalated as needed. Other important elements that an organisation may include in an IR plan are:

- documentation;
- preservation of forensic evidence, including chain of custody;
- preservation of privilege;
- responding to a ransom demand; and
- third-party incidents.

An organisation may also decide to reference or include information regarding other related processes, such as business continuity and disaster recovery plans.

Maintaining and testing an incident response plan

After developing an IR plan, organisations should undertake regular reviews and testing to help ensure that the plan remains current and effective over time. Even the best plans require updating and testing to remain effective.[13]

Reviews

Regular review and revision of an IR plan is vital to maintaining it as a functional, practical document. Generally, a plan should be reviewed at least annually to ensure critical information is current when an incident occurs. Among other elements, the review should confirm that personnel and contact information is up to date, legal and regulatory developments are considered and incorporated as appropriate and any organisational changes, such as internal restructuring or changes to group functions, are reflected. Significant personnel and legal, regulatory and organisational developments should themselves trigger review of the plan, without waiting for any scheduled maintenance review. Any modifications to the plan that result from a review should be recorded in a change log that includes the date of the modification and the person responsible for the changes.

In addition to scheduled reviews, it is good practice to review the plan after an incident requiring plan activation. Depending on the organisation's overall approach to incident severity classification and plan activation, that review may reasonably occur after incidents of a certain severity. This post-incident plan review can be incorporated as part of the IR process described above, that is, as part of the post-incident activity. A post-incident review will have a somewhat different purpose than a scheduled review

and will focus on incorporating lessons learned from the IR process. For example, if during the IR process a different member of the CIRT from the member designated in the plan handled communications to external parties, and that worked well, the organisation might decide to change that assignment in the plan.

Maintenance and incident reviews are relatively straightforward processes that prevent the plan from becoming outdated or reliant on ineffective responses when it needs to be used.

Testing

Another important step for maintaining the effectiveness of an IR plan is to simulate an incident through tabletops and other exercises. Because the organisation controls the scenario, the most productive and useful simulations will target areas of potential weakness so that the organisation can grow, learn and allocate resources to address such weaknesses. Tabletop exercises also allow an organisation to assess whether the CIRT and other personnel understand their roles, responsibilities and protocols during an incident, and, if not, whether additional training is required.

Tabletop exercises can also target different functions and areas of an organisation, and likewise test different aspects of a plan. For example, tabletop exercises for management or board members can occur in a boardroom with the tabletop coordinator providing different organisation-specific scenarios, and the participants discussing how to respond at a high level with a heavier focus on issues such as disclosure. By contrast, technical tabletop exercises designed to test the members of the CIRT can be more involved, for example, with a controlled 'attack' on the organisation by the tabletop facilitator requiring the participants to defend and eradicate the 'threat'. Some organisations also run tabletop exercises with third parties, testing the response capabilities of key vendors and suppliers and their ability to coordinate during an incident. All types of tabletop exercises are important measures to prepare an organisation for an actual incident and ensure that the IR plan is effective and practical.

Conclusion

Every organisation should invest in a cybersecurity incident response plan suited to its unique needs so that when an incident occurs it is well positioned to protect its customers, business, employees and data. A thoughtful approach during the development process will yield a plan that encompasses key elements and will be useful to the organisation in a crisis.

Endnotes

[1] The value of an IR plan is supported by data: IBM's 2023 'Cost of a Data Breach Report' determined that, on average, organisations with high levels of IR planning and testing incurred $1.49 million less in data breach costs and resolved incidents 54 days faster, and that IR preparation and testing were among the top-three cost mitigators. IBM, 'Cost of a Data Breach Report' 66 (2023).

[2] This chapter focuses on the key elements for IR plans aimed at guiding the IR response process. We do not address the various types of detailed IR playbooks used by information technology and information security professionals to respond to the technical aspects of an incident.

[3] NIST defines an 'event' as 'any observable occurrence that involves computing assets'; 'adverse events' as 'any events associated with a negative consequence regardless of cause'; and 'incident' as 'an occurrence that actually or imminently jeopardises, without lawful authority, the integrity, confidentiality, or availability of information or an information system; or constitutes a violation or imminent threat of violation of law, security policies, security procedures, or acceptable use policies'. Nat'l Inst. of Standards & Tech., U.S. Dep't of Comm., NIST SP 800-61r3, Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile 2 (Apr. 2025), available at **https://doi.org/10.6028/NIST.SP.800-61r3**.

[4] See 17 C.F.R. § 229.106(a) (defining a '[c]ybersecurity incident' as 'an unauthorized occurrence, or series of related unauthorized occurrences, on or conducted through a registrant's information systems that jeopardizes the confidentiality, integrity, or availability of a registrant's information systems or any information residing therein').

[5] See, e.g., Md. Code Ann., Com. Law § 14-3501(e)(1); Del. Code Ann. tit. 6, § 12B-101(7).

[6] See, e.g., Cal. Civ. Code §§ 1798.80(e), 1798.82.

[7] 2016 O.J. (L 119) 33. As to other jurisdictions, Canada's Personal Information Protection and Electronic Documents Act provides that 'personal information means information about an identifiable individual'. Personal Information Protection and Electronic Documents Act, S.C. 2000, c 5, § 2(1) (Can.). Similarly, South Africa's Protection of Personal Information Act defines 'personal information' as 'information relating to an identifiable, living, natural person' or any 'identifiable, existing juristic person'. Protection of Personal Information Act of 2013 § 1 (S. Afr.).

[8] See, e.g., Alaska Stat. § 45.48.090(1); Ind. Code § 24-4.9-2-2(a).

[9] Compare Fla. Stat. § 501.171(1)(a) (defining 'breach' as the 'unauthorized access of data in electronic form containing personal information') with Kan. Stat. Ann. § 50-7a01(h) (defining 'breach' as 'unauthorized access and acquisition of unencrypted or unredacted computerized data that compromises the security, confidentiality or integrity of personal information').

[10] GDPR art. 4(12). Under Singapore's data privacy statute, a data breach broadly includes any 'unauthorized access, collection, use, disclosure, copying, modification or disposal of personal data', regardless of whether any harm or risk of harm was caused by the breach. Personal Data Protection (Amendment) Act 2020 § 26A.

[11] 45 C.F.R. § 164.400-414.

[12] See, e.g., NIST SP 800-61r3, supra note __, at 6.

[13] In the wake of an incident, regulators and plaintiffs may also request to review an organisation's IR plan. If that review occurs, it is significantly better if the plan is updated and reflects positively on the organisation's ongoing compliance efforts and pursuit of continuous improvement.

Evan Norris

Partner
Cravath, Swaine & Moore LLP
enorris@cravath.com

Dean M Nickles

Of counsel
Cravath, Swaine & Moore LLP
dnickles@cravath.com