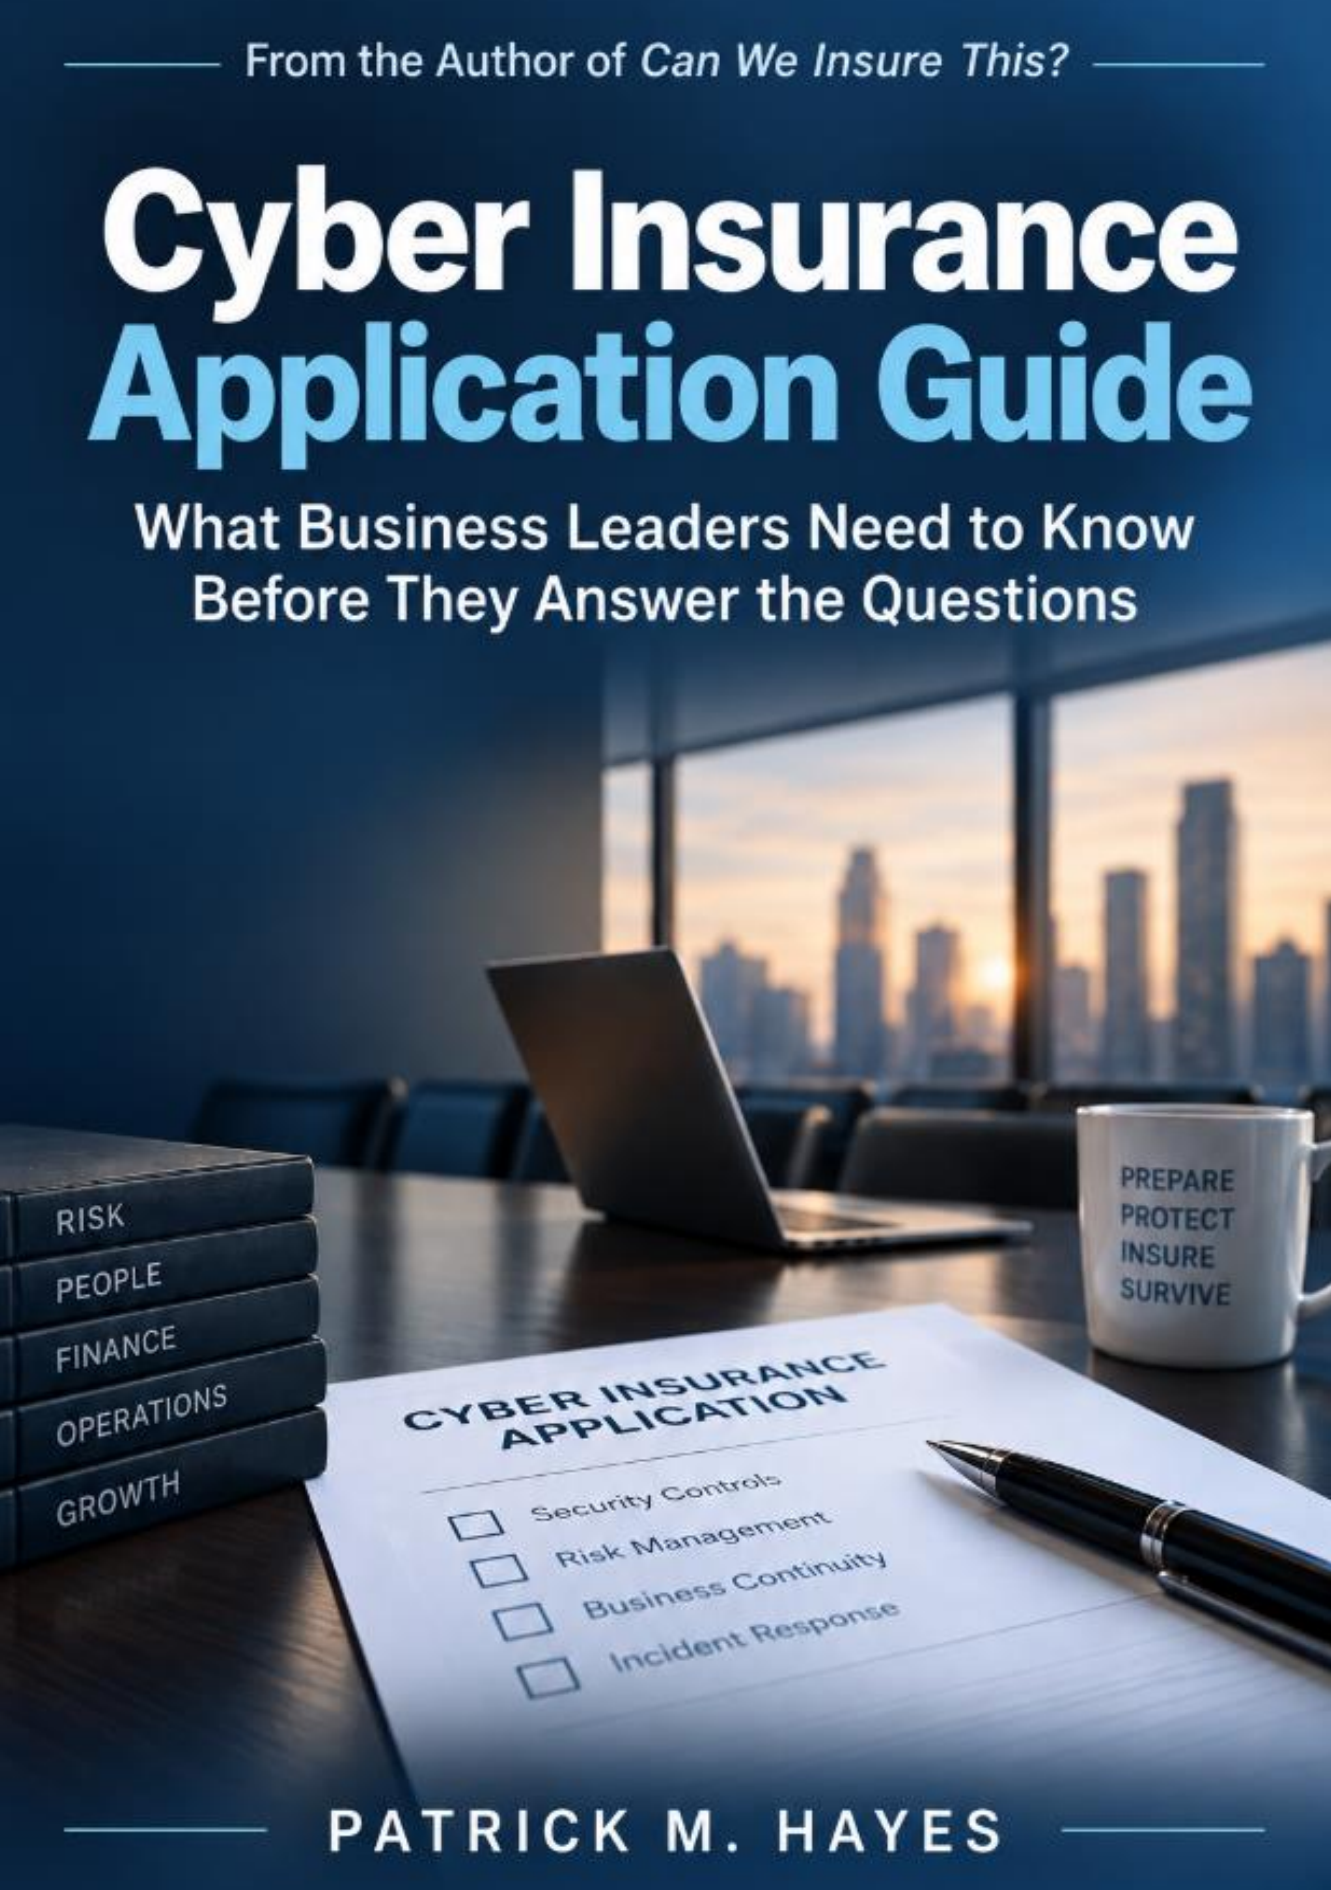
From the Author of *Can We Insure This?*

Cyber Insurance Application Guide

What Business Leaders Need to Know Before They Answer the Questions



PATRICK M. HAYES

© All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations used in critical reviews or scholarly works with proper attribution.

Table of Contents

Understanding the Cyber Insurance Application and Renewal Process 5
- Why Insurers Are Asking These Questions 5
- Why Cyber Insurance Matters 6
- What Information the Insurer Wants 6
- Knowing the Answer Is Different from Assuming the Answer 7
- Collecting Evidence Before Answering 7
- Finding the Right Person Inside the Business 8
- Creating an Underwriting Evidence File 8
- What to Do When the Answer Is Not Yes 9
- Making Renewal a Business Process 9

Access Control 11
- Do you have Single Sign-On (SSO) implemented in your organization? 11
- Do you enforce multi-factor authentication for access to systems and applications? 12
- Do you restrict internet access to business-critical technologies? 12
- Do you enforce strong passwords? 13
- Do you rotate user passwords periodically? 13
- Do you limit access to critical information based on need-to-know and least privilege? 14
- What the Executive Should Take Away from Access Control 15

Human Risk 16
- Do you conduct periodic security awareness training for all employees? 16
- Do you conduct periodic phishing simulations or exercises? 17
- What the Executive Should Take Away from Human Risk 18

Email Security 19
- Have you implemented an Email Security solution to filter spam or malicious emails? 19
- What the Executive Should Take Away from Email Security 20

Endpoint Security 22
- Which basic endpoint security controls have you implemented across your devices? 22
- Have you implemented an Endpoint Detection and Response (EDR) solution on your devices? 23
- Do you use a Managed Detection and Response (MDR) service today? 24
- What the Executive Should Take Away from Endpoint Security 25

Network Security 26
- Do you have a firewall in place for your network? 26

Is the guest network separate from the corporate or production network? 27
What the Executive Should Take Away from Network Security 28

Patch Management 29
Do you have a process to patch or address critical vulnerabilities? 29
Do you have a Managed Service Provider (MSP) or Managed Security Services Provider (MSSP)? 30
What the Executive Should Take Away from Patch Management 31

Data Backup 33
Do you have a backup process for your data? 33
Do you have a plan to restore business functionality and data in case of a disaster? 34
What the Executive Should Take Away from Data Backup 35

Incident Response 37
Do you have an incident response plan? 37
Are incidents logged and tracked? 38
Is there a process for reporting incidents to appropriate authorities? 39
What the Executive Should Take Away from Incident Response 40

Personnel Security 41
Are background checks conducted for employees? 41
Is there a process for handling employee terminations? 42
What the Executive Should Take Away from Personnel Security 42

Can We Insure This? 44

Understanding the Cyber Insurance Application and Renewal Process

Cyber insurance applications have changed quite a bit over the past few years. What used to be a relatively simple questionnaire about firewalls, antivirus software, and a handful of basic security practices has become a much more detailed look at how a business actually operates. That change makes sense when you consider what insurers are trying to understand. They are not really trying to grade your cybersecurity program. They are trying to determine how likely your organization is to have a cyber event, what that event could cost, and whether the business has the ability to limit the damage when something eventually goes wrong.

The cyber insurance application should not simply be forwarded to IT with a request to answer the technical questions. The answers provided on the application become representations about the business that an insurer may rely upon when deciding whether to provide coverage and under what terms. Those answers can influence premiums, deductibles, coverage limits, exclusions, and other policy conditions. They can also become important after a loss when the insurer begins looking at what happened and comparing the actual environment with what was represented during underwriting.

Why Insurers Are Asking These Questions

An insurer is ultimately trying to understand the financial risk it is agreeing to take. A cyber incident can create expenses far beyond the cost of fixing a computer or restoring a server. The business may lose revenue while systems are unavailable. Employees may be unable to work. Customers may be affected. Outside forensic investigators and attorneys may be needed. Data may need to be restored, regulators may become involved, and contractual obligations may create additional costs. The insurer needs enough information to estimate how likely those situations are and how severe the financial impact could become.

This is why applications have become more detailed. Insurers have learned from claims which conditions tend to make losses worse. Weak access controls can make account compromise easier. Poorly managed endpoints can allow an attacker to spread through the organization. Untested backups can turn an otherwise manageable incident into a prolonged outage. An incident response plan that exists only as a document may provide little value when nobody knows what to do during an actual event. The questions on the application are often trying to identify those conditions before the insurer agrees to take on the risk.

Why Cyber Insurance Matters

Cyber insurance has practical value because a serious cyber event can create costs that most businesses are not prepared to absorb on their own. A ransomware attack, data breach, system outage, fraudulent payment, or compromised email account can quickly lead to forensic expenses, legal fees, recovery costs, customer notification, business interruption, and lost revenue. The policy is there to help finance covered losses when the impact moves beyond something the company can handle as a normal operating expense.

The policy can also give the business access to specialized help at the point when it is most needed. Depending on the coverage and carrier, that may include breach counsel, forensic investigators, incident response firms, notification providers, negotiators, and other specialists that many companies would not already have under contract. That support matters because most businesses do not deal with major cyber incidents often enough to build those capabilities internally, and trying to find the right experts in the middle of an event can waste valuable time.

Cyber insurance does not replace good security, recovery planning, or business continuity. It helps address the financial consequences that remain after the company has done what it reasonably can to reduce the risk. That is why the application and renewal process matters so much. The insurer is trying to understand the business it may be asked to protect, while the business should be making sure the policy it buys actually lines up with the losses it would struggle to absorb on its own.

What Information the Insurer Wants

Most cyber insurance applications cover the same general areas even though the wording and level of detail can vary considerably between carriers. The application you are completing may ask about access control, human risk, email security, endpoint security, network security, patch management, data backup, incident response, and personnel security. Each category gives the insurer another piece of the overall picture of how your organization manages cyber risk.

Access control questions help the insurer understand who can get into important systems and what prevents someone from gaining access using stolen credentials. Human risk questions look at how employees are prepared to recognize threats and avoid actions that could expose the organization. Email security receives considerable attention because email continues to be involved in many fraud and account-compromise scenarios. Endpoint and network questions help the insurer understand how devices and systems are protected and whether suspicious activity can be identified before it becomes a larger problem.

Patch management questions usually focus on whether known vulnerabilities are being addressed before attackers can exploit them. Backup questions are particularly important because the ability to restore systems and information can have a major effect on the size of a business interruption loss. Incident response questions help the insurer understand whether the organization is prepared to manage an event once it occurs. Personnel security extends the discussion to how employees and contractors are managed when they join the organization, change responsibilities, or leave.

None of these areas should be viewed in isolation. The insurer is trying to develop a reasonable picture of what would happen if the business were attacked tomorrow. The real issue is whether an incident would remain manageable or develop into a significant financial loss.

Knowing the Answer Is Different from Assuming the Answer

This is where many businesses get themselves into trouble. Someone completing the application asks IT whether the company uses multifactor authentication and receives a quick "yes." The box gets checked and everyone moves on. The problem is that the insurer may not have asked whether the company uses MFA somewhere. It may have asked whether MFA is required for remote access, administrative accounts, email, cloud applications, or some other specific environment.

The same thing happens with backups. A business may confidently answer that backups are performed every day without knowing whether those backups are separated from the production environment or whether anyone has successfully restored critical systems from them recently. A company may say that endpoint protection is deployed without knowing whether every endpoint is actually covered. These differences can seem technical during the application process, but they can become very important after a claim.

The person completing the application does not need to become a cybersecurity expert. They do, however, need to understand exactly what is being represented to the insurer and make sure the answer has been verified by someone who actually knows.

Collecting Evidence Before Answering

The easiest way to improve the application process is to stop treating it as a questionnaire and start treating it as a verification exercise. When the insurer asks an important question about a security control, the organization should be able to identify some reasonable evidence supporting the answer.

That evidence will look different depending on the question. An MFA question might be supported by a configuration report from Microsoft 365, a VPN platform, or another identity system. An endpoint security answer might be supported by a report showing the devices currently enrolled in the company's endpoint protection platform. Backup questions could be supported by backup reports and records from restoration tests. Training questions might be supported by completion records from the organization's security awareness platform. Patch management could involve vulnerability reports or records from the system used to deploy updates.

The point is not to create a huge collection of screenshots just to satisfy an insurer. The evidence should give the business reasonable confidence that the answer being submitted is actually correct. If nobody can find evidence supporting an answer, that is worth investigating before checking the box.

Finding the Right Person Inside the Business

One reason cyber insurance applications become frustrating is that the person responsible for completing the application rarely has all of the answers. That is normal. The information needed to complete the application may exist across several parts of the organization and, increasingly, with outside service providers.

IT will usually have much of the information involving accounts, devices, networks, patching, and backups. A security team or managed security provider may have information about monitoring, endpoint protection, vulnerabilities, and incident response. Human resources may need to verify employee training, onboarding, background screening, and termination practices. Finance may need to answer questions involving revenue, payment processes, business interruption, and financial controls. Legal or compliance teams may need to help identify regulated information, contractual obligations, previous incidents, and notification requirements.

Smaller businesses may depend heavily on an MSP or another technology provider for many of these answers. That does not eliminate the need for verification. If a provider says it manages backups, endpoint security, or MFA for the business, the organization should understand what is actually being provided and should be able to obtain evidence that supports the answer being given to the insurer.

Creating an Underwriting Evidence File

A relatively simple improvement is to maintain an underwriting evidence file alongside the insurance application. For each material question, the business

should record the answer that was submitted, who verified the answer, what evidence was used, and when that verification occurred. The evidence itself can then be retained with the application and policy documents.

This makes the renewal process much easier because the business is not starting from scratch every year or trying to remember why someone answered a question a certain way twelve months earlier. It also gives leadership a much better way to identify changes. A control that existed at the previous renewal may have disappeared because a system was replaced. A new acquisition may have introduced technology that is not covered by the same controls. An outside provider may have changed. The evidence provides a baseline that can be reviewed instead of relying on assumptions.

The process can also uncover problems that have nothing to do with buying insurance. If the business discovers that some users are not protected by MFA, that backups have not been tested, or that terminated employees still have active accounts, the application has identified an operational issue worth fixing. That is one of the more useful aspects of a well-run renewal process.

What to Do When the Answer Is Not Yes

There can be pressure during an insurance application to turn every answer into a "yes," particularly when everyone knows that certain answers may affect the premium or whether coverage is offered. That is the wrong objective. The goal should be to provide an accurate answer and understand what the insurer considers important.

Sometimes the correct answer will be "no." In other situations, the control may exist but only cover part of the environment. That should lead to a conversation with the broker or insurer about exactly what the question means and whether additional information can be provided. It may also identify something the organization can reasonably correct before the application is finalized.

Finding a weakness during underwriting is usually much better than discovering it during a claim.

Making Renewal a Business Process

Cyber insurance renewal works better when it begins before the application arrives. The organization should review the previous year's application, confirm that the important answers are still accurate, update the supporting evidence, and identify significant changes that have occurred in the business. New systems, acquisitions, cloud migrations, vendors, remote access, AI adoption, and changes

in how the company handles information can all affect the answers that were previously provided.

For executives and financial leaders, the value of this process goes beyond obtaining an insurance policy. The application provides a structured opportunity to compare what leadership believes about the organization with what can actually be demonstrated. When those two things do not match, the business has learned something important.

The sections that follow will walk through each area of the application in detail. For every question, the focus will be on understanding what the insurer is really trying to determine, why the answer matters to underwriting, where the information can be found inside the organization, and what evidence should be collected before an answer is provided.

Access Control

Access control is one of the first areas insurers look at because many serious cyber incidents begin with someone gaining access they should not have. Sometimes an attacker steals an employee's password. In other cases, an administrative account is compromised or a system that should not have been accessible from the internet is discovered. The insurer is trying to understand how difficult it would be for someone to get into the organization and, once they are in, how much access that person could have.

For a business leader, the important thing to understand is that these questions are not simply asking whether certain technologies exist somewhere in the company. Words such as all users, all technologies, enforce, and only matter. If a control protects 95 percent of the organization, an application asking whether it protects all users may not support a "Yes" answer. This is where evidence becomes particularly important.

Do you have Single Sign-On (SSO) implemented in your organization?

Single Sign-On allows employees to use a central identity to access multiple business applications rather than maintaining separate usernames and passwords for every system. Microsoft Entra ID, formerly Azure AD, Okta, Google, and Ping Identity are common examples. From an insurer's perspective, SSO matters because centralized access gives the business greater control over who can access its systems. When an employee leaves, for example, disabling the central account can remove access to many connected applications instead of depending on someone to remember every individual account the employee had.

The important part of this question is the difference between having SSO and having it broadly implemented. A company might use Microsoft Entra ID to access Microsoft 365 while employees still use separate passwords for accounting software, customer systems, cloud platforms, or other important applications. In that situation, saying that SSO is rolled out to all users may not tell the entire story. The person completing the application should find out which applications are actually connected to SSO and whether important exceptions exist.

The best source for this answer will normally be the internal IT team, identity administrator, or MSP responsible for managing user accounts. Evidence could include an identity platform configuration report, a list of applications connected to the SSO platform, or administrative screenshots showing the configuration. If the insurer asks which SSO provider is being used, that should be easy to verify through the same source.

Do you enforce multi-factor authentication for access to systems and applications?

MFA has become particularly important to cyber insurers because a stolen password by itself should not be enough to give an attacker access to an important system. The second factor creates another barrier. For insurers, however, the important word in this question is usually enforce. Having MFA available to employees is not the same as requiring them to use it.

This question also requires some care because MFA may exist in one part of the company but not another. Microsoft 365 may require MFA while a VPN, firewall, server, remote desktop connection, cloud application, or administrative account does not. When the application asks whether MFA applies across all users, the business should verify the scope before answering. The follow-up question about email, remote access, internet-facing systems, servers, and user devices is intended to uncover exactly this kind of difference.

The insurer is also asking what happens after failed MFA attempts because repeated failures can indicate that someone is attempting to access an account. Simply having MFA is useful, but the insurer also wants to know whether suspicious authentication activity is detected and acted upon.

IT, the identity administrator, security team, or MSP should provide the answer. Useful evidence includes MFA enforcement reports, conditional access policies, identity-provider configuration, VPN configuration, firewall settings, and reports showing MFA enrollment and coverage. The evidence should demonstrate enforcement rather than simply showing that the MFA feature has been enabled.

Do you restrict internet access to business-critical technologies?

This question can be confusing because it sounds as though the insurer is asking whether employees are allowed to browse the internet. That is not really the issue. The insurer wants to know whether important technology can be reached directly from the internet when there is no business reason for it to be exposed.

An internet-facing administrative portal, remote management service, server, firewall interface, or remote desktop service can provide an attacker with a place to begin trying to enter the organization. The insurer therefore wants to understand whether those systems have been identified and whether access is limited through firewalls, VPNs, access restrictions, or other controls.

The follow-up questions give a better indication of what the insurer is looking for. It wants to know whether internet-facing technologies have only the access they actually require, whether remote management tools are protected, and whether powerful administrative access is controlled. A business should not answer "Yes, across all technologies" simply because it has a firewall. Someone needs to confirm what is actually exposed to the internet.

The internal IT team, network administrator, security provider, or MSP will normally have this information. Evidence could include firewall configurations, external vulnerability scan results, network diagrams, remote access configurations, VPN settings, or documentation showing which systems and services are intentionally exposed to the internet.

Do you enforce strong passwords?

This sounds like one of the easiest questions on the application, but it is easy to answer based on policy rather than reality. A written policy stating that employees must use strong passwords does not necessarily mean the requirement is technically enforced.

The insurer is trying to determine how easily an attacker could guess, steal, or reuse credentials to gain access to the organization. The business should determine what password requirements are actually configured, whether those requirements apply consistently, and whether older systems or applications operate under different rules.

The answer will usually come from IT, the identity administrator, or the MSP. Evidence might include Microsoft Entra ID or Active Directory password settings, identity-provider configuration, application security settings, or other system-generated documentation showing the password requirements being enforced. A copy of the company's password policy can support the answer, but it is stronger when accompanied by evidence showing that the technology actually enforces the policy.

Do you rotate user passwords periodically?

This question deserves some explanation because cybersecurity guidance around password rotation has changed. Periodically forcing every employee to change a password was once considered standard practice. More recent guidance has moved away from requiring arbitrary password changes when there is no indication that a password has been compromised. Strong passwords, MFA, compromised-password detection, and changing credentials when compromise is suspected can provide a better approach.

That does not make the insurer's question irrelevant. The application is asking what the organization actually does today. If passwords are changed every 90 days, that should be verified. If the organization intentionally does not require periodic changes because it follows a different password-management approach, the answer should reflect that rather than selecting a more favorable-looking response.

This is also a good example of why an executive should not try to interpret the desired answer and then ask IT to confirm it. The first question should simply be, "What do we actually do?" The broker can then help determine whether additional explanation should be provided to the insurer.

IT or the identity administrator should verify the answer. Password configuration settings, identity policies, Active Directory or Entra ID settings, and password-management platform reports can provide evidence. If the organization follows a documented password standard that does not require routine rotation, retaining that documentation can also help explain the answer.

Do you limit access to critical information based on need-to-know and least privilege?

This question is really about how much damage one compromised account could cause. Employees generally do not need access to every system or every piece of company information. Administrative accounts are particularly important because they can often change configurations, create users, disable security controls, access sensitive information, or make other significant changes.

The insurer wants to know whether powerful access is limited to people who actually need it. It is also looking at whether default accounts and default configurations are addressed when systems are installed. Default administrator accounts and unnecessary privileges can give attackers opportunities that should have been removed during normal system setup.

For executives, there is a useful distinction here. Giving administrative privileges only to a subset of employees does not automatically mean the company has implemented least privilege throughout the organization. The real question is whether people's access reasonably reflects what they need to do their jobs and whether that access changes when their responsibilities change.

IT, system owners, security teams, application administrators, and MSPs may all hold pieces of this answer. Human resources may also be involved because job changes and terminations should trigger access changes. Evidence can include lists of administrative accounts, role and group membership reports, access

reviews, privileged access management reports, onboarding and termination procedures, and records showing that unnecessary or default accounts have been disabled.

What the Executive Should Take Away from Access Control

Access control questions are a good example of why completing a cyber insurance application requires more than knowing that the company has purchased the right security tools. An organization can own an identity platform, use MFA, maintain a firewall, and have a password policy while still having significant gaps in how those controls are applied.

When reviewing these answers, the executive or financial leader does not need to personally validate technical configurations. They should, however, know who verified the answer and what evidence supports it. If the answer is "Yes, across all users" or "Yes, across all technologies," someone should be able to demonstrate that coverage. If exceptions exist, they should be understood before the application is submitted.

The objective is not to get six perfect answers. It is to give the insurer accurate information about the organization as it exists today and make sure the business understands what it is representing when someone signs the application.

Human Risk

Human risk questions are intended to help the insurer understand how prepared employees are to recognize and respond to the types of attacks that specifically target people. This matters because attackers often do not need to defeat sophisticated security technology if they can convince an employee to give them access, open a malicious attachment, approve a fraudulent payment, or provide information that helps them get further into the organization.

For business leaders, these questions are relatively straightforward, but there is an important distinction between having a training program and being able to demonstrate that employees actually participate in it. An annual presentation or an email reminding employees to be careful is different from a managed program that records participation and provides evidence that training occurred. The insurer is looking for evidence that the organization is actively addressing the risks created by normal human behavior.

Do you conduct periodic security awareness training for all employees?

Security awareness training teaches employees how to recognize situations that could put the company at risk. That can include phishing emails, suspicious attachments, fraudulent requests, password theft, social engineering, unsafe use of company information, and other situations employees are likely to encounter while doing their jobs.

The insurer asks this question because employees are frequently the target of an attack. A finance employee might receive what appears to be an urgent payment request from an executive. Someone in human resources could receive a request for employee information. An employee could be directed to a fake Microsoft login page and unknowingly provide their credentials. The insurer wants to know whether employees have been given enough training to recognize these situations before they become insurance claims.

The wording "for all employees" is important. A company may have a security awareness program while still having employees who have not completed the training. New employees may have joined after the last training cycle, some employees may have ignored the training request, or contractors and temporary workers may fall outside the program. Before answering that everyone receives annual training, the organization should confirm actual participation rather than relying on the fact that training was made available.

The answer will usually come from human resources, IT, the security team, or the MSP or security provider managing the program. In smaller organizations, the

office manager or another person responsible for employee training may have the information. If a platform such as KnowBe4, Proofpoint, SoSafe, NINJIO, or another service is being used, the platform itself will normally contain much of the evidence needed.

Good evidence would include a training completion report showing the employees assigned training, completion dates, and completion status. The business could also retain the training schedule and records showing how new employees are added to the program. Simply providing an invoice showing that the company purchased security awareness software does not prove that employees actually completed the training.

The optional question about the training provider and contract expiration is also worth answering accurately. The insurer is trying to understand what program exists today and whether it is likely to remain in place. The contract information will normally be available from IT, procurement, finance, the MSP, or whoever purchased the service.

Do you conduct periodic phishing simulations or exercises?

Phishing simulations go a step beyond awareness training by testing how employees respond when they receive something that looks like a real phishing attempt. Employees may receive a simulated malicious email containing a link, attachment, login request, or another scenario designed to determine whether they recognize the warning signs.

The insurer cares about this because completing training does not necessarily mean someone will recognize an attack when it appears in their inbox six months later. Phishing simulations give the organization some evidence of how employees actually respond. They can also identify people or parts of the business that may need additional training.

The frequency matters here. If the application asks whether simulations are conducted monthly, quarterly, or annually, the answer should reflect what actually occurred during the period rather than what the company's policy says should occur. A company that intended to conduct quarterly testing but only completed one exercise during the year should not describe the program as quarterly.

The results themselves can also be useful. A good phishing program should allow the organization to see how many employees received the simulation, how many interacted with it, whether credentials were entered, whether employees reported the suspicious message, and whether additional training was assigned when necessary. The purpose is not to embarrass employees who make mistakes. It is

to understand whether the organization is becoming more resistant to the types of attacks employees will eventually encounter.

The security team, IT department, MSP, or security awareness provider will normally have this information. Evidence can include phishing campaign reports showing when exercises occurred, how many employees participated, and the results of those exercises. If additional training was assigned after a failed simulation, those records can provide additional support that the program is actively managed.

As with the previous question, naming the phishing simulation platform is useful but does not by itself prove that simulations are being conducted. A subscription to a product and an operational phishing program are not necessarily the same thing.

What the Executive Should Take Away from Human Risk

These questions should give leadership a fairly quick way to determine whether the company has a functioning security awareness program or simply believes that it does. If someone says that every employee receives annual security training, there should be a report showing that. If the company says phishing exercises happen quarterly, there should be records of those exercises.

This is also an area where the evidence is usually easy to obtain. Most modern awareness and phishing platforms already produce the reports an insurer would need. The challenge is making sure someone actually reviews those reports and follows up when employees do not complete training or repeatedly struggle with simulations.

From an insurance perspective, the company should be able to show that the program existed and was operating at the time the application was completed. From a business perspective, the more important question is whether employees are actually becoming better at recognizing the attempts that could lead to a financial loss.

Email Security

Email security receives a lot of attention from cyber insurers because email remains one of the easiest ways to reach someone inside a business. An attacker does not necessarily need to break through a firewall if they can send an employee a convincing message, steal their Microsoft 365 credentials, get them to open a malicious attachment, or convince someone in finance to change payment information. The resulting loss can involve ransomware, data theft, account compromise, fraudulent payments, and business interruption.

For business leaders, the important point is that having Microsoft 365 or Google Workspace does not automatically answer the insurance question. Those services include security capabilities, but the protections actually being used depend on licensing, configuration, additional security products, and how the environment is managed. The organization needs to understand what is protecting its email today rather than assuming that email security came with the email service.

Have you implemented an Email Security solution to filter spam or malicious emails?

The insurer is asking whether something is actively examining incoming and, in some cases, outgoing email before a potentially dangerous message reaches an employee. Most organizations have some level of basic spam filtering, but the insurer is interested in protections that can identify phishing attempts, malicious links, dangerous attachments, impersonation attempts, malware, and other suspicious email activity.

This matters because the email service and the email security solution may not be the same thing. A company may use Microsoft Exchange Online for email while relying on Microsoft Defender for Office 365, Proofpoint, Abnormal Security, Material Security, or another product for additional protection. Another company may rely entirely on the protections included with its Microsoft or Google environment. Neither situation should be assumed. Someone needs to verify what is actually licensed, configured, and operating.

The follow-up question asking which email service the organization uses establishes where business email is hosted. For most companies, the answer will be Microsoft Exchange Online, Google Workspace, an on-premise Exchange environment, or another hosted service. This information should be easy for IT or the organization's technology provider to verify.

The question about specific email security controls requires more investigation. Anti-spam and anti-phishing protections are intended to prevent unwanted or

fraudulent messages from reaching employees. Malware and ransomware protection examines attachments, links, and other email content for malicious activity. Data Loss Prevention, commonly called DLP, addresses a different problem by helping prevent sensitive company information from being sent somewhere it should not go. A company should only select these controls when they are actually configured and in use.

The insurer also asks which email security product is being used because products provide different capabilities and levels of protection. This does not mean that simply owning one of the listed products produces a better answer. A security product that was purchased but poorly configured may provide considerably less protection than leadership believes. The insurer is trying to understand what actually sits between an attacker and the company's employees.

IT, the email administrator, security team, or MSP will normally be the best source for these answers. If email security is managed by an outside provider, that provider should be able to explain exactly which product is being used and which protections it manages on the company's behalf.

Evidence could include administrative screenshots or reports showing that the email security service is active, licensing information showing which security capabilities are available, configuration reports showing enabled protections, and reports showing that malicious or suspicious messages are actually being detected or blocked. If DLP is selected, there should be evidence showing that DLP policies are configured and applied rather than simply showing that the technology supports the feature.

There is also a useful distinction between evidence that a product was purchased and evidence that a control is operating. An invoice for Microsoft Defender for Office 365 or Proofpoint confirms that the company bought something. It does not necessarily prove how it is configured or whether it protects every employee covered by the application. For underwriting purposes, configuration and operational reports provide much stronger support for the answer.

What the Executive Should Take Away from Email Security

The executive reviewing this section does not need to understand the technical details of email filtering. The important question is whether someone can clearly explain what protects the company's email environment and provide evidence that those protections are actually operating.

This is another area where assumptions can create problems during an insurance application. Someone may know that the company uses Microsoft 365 and

reasonably assume that Microsoft handles email security. That may be partly true, but it does not establish which protections are enabled or whether the organization has purchased the licensing required for some of the capabilities it believes it has.

The person approving the insurance application should be able to identify the company's email platform, the security solution protecting it, and the controls that are actually enabled. The supporting evidence should come from the people managing the environment rather than from what the business believes was included when the service was purchased.

Endpoint Security

Endpoint security is where the insurer starts looking more closely at the computers and other devices people actually use to operate the business. Laptops, desktops, servers, and other connected devices are attractive targets because they provide access to applications, company information, credentials, and often the rest of the technology environment. If an attacker compromises one of these devices, the insurer wants to understand what prevents that initial compromise from becoming a much larger event.

For business leaders, this section can become confusing because several different security technologies are involved and their names sound similar. Antivirus, endpoint protection, EDR, and MDR are not interchangeable. It is possible for a company to have one without having the others. The application is trying to understand both what technology is installed and what happens when that technology detects something suspicious.

Which basic endpoint security controls have you implemented across your devices?

This question establishes the basic protections that exist on company devices. The insurer is asking about several different controls because each addresses a different type of risk. A firewall controls network communications to and from the device. Disk encryption protects information stored on a laptop or other device if it is lost or stolen. Antivirus or an endpoint protection platform attempts to prevent malicious software from running. A VPN can protect connections when employees access company resources remotely. Device trust can be used to determine whether a device meets the company's requirements before allowing it to connect to important systems.

The phrase "across your devices" deserves attention. A company may have these protections on its standard corporate laptops but not on older computers, servers, employee-owned devices, remote systems, or devices used at smaller locations. The person completing the application should understand what population of devices is covered before selecting a control.

This is also why simply knowing that the company owns endpoint security software is not enough. If the organization believes it has 500 computers but the endpoint management platform shows only 430 protected devices, someone should understand what accounts for the difference before answering the insurance question.

IT, the endpoint administrator, security team, or MSP will normally provide this information. Evidence could include device inventory reports, endpoint management reports, encryption status reports, firewall configuration, VPN deployment records, and security platform reports showing which devices are actively protected. Ideally, the organization should be able to compare its device inventory with the security platform and explain any meaningful difference.

Have you implemented an Endpoint Detection and Response (EDR) solution on your devices?

EDR goes beyond traditional antivirus. Antivirus is primarily intended to prevent known malicious software from running. EDR watches activity occurring on a device and looks for behavior that could indicate an attack, including activity that may not match a known piece of malware. Depending on the product and configuration, EDR can also provide the ability to investigate suspicious activity, isolate a compromised computer, and take other actions to contain an incident.

Insurers care about EDR because preventing every attack is unrealistic. What matters after something gets through is whether the organization can detect it before the attacker has enough time to move through the environment, steal information, disable systems, or deploy ransomware.

The question should not be answered based only on whether the company purchased an EDR product. The organization needs to understand whether the EDR software is actually deployed to the devices it is supposed to protect and whether those devices are actively reporting into the platform. A deployment that covers most laptops but leaves important servers unprotected may need to be described differently depending on how the insurer's question is written.

The follow-up question asking which EDR provider is used helps the insurer understand what capabilities are available. The product name is useful, but the more important issue is how extensively it has been deployed and how it is being operated.

IT, the security team, endpoint administrator, SOC, or MSP should be able to verify this information. Strong evidence would include an EDR console report showing active devices, deployment status, policy assignment, and recent communications from those devices. Comparing that information with the organization's device and server inventory provides much stronger evidence than an invoice showing that EDR licenses were purchased.

Do you use a Managed Detection and Response (MDR) service today?

This is where an important distinction needs to be made. EDR is primarily the technology watching the endpoints. MDR is a service involving people who monitor security information, investigate suspicious activity, and respond when something requires attention.

A company can therefore have EDR without having MDR. The EDR platform may generate alerts while someone inside the organization is expected to investigate them. For a small or midsized business without a dedicated security team, that can create a problem. An alert that appears at 2:00 Sunday morning does little good if nobody looks at it until Monday.

The insurer asks about MDR because it wants to know what happens after the security technology identifies something suspicious. A managed service may provide monitoring outside normal business hours and access to security specialists who can investigate an event. Depending on the agreement, the provider may also be authorized to isolate a compromised device or take other containment actions.

Business leaders should be particularly careful about assuming that an MSP automatically provides MDR. An MSP may manage computers, Microsoft 365, backups, and other technology without providing continuous security monitoring. Likewise, a company may purchase an EDR product through its MSP without purchasing a service that actively monitors the alerts generated by that product.

The person reviewing the insurance application should ask the provider what happens when the EDR platform generates a serious alert at night or over a weekend. The answer will usually make it fairly clear whether the organization has a managed detection capability or simply owns detection technology.

Evidence should come from the MDR provider, security team, or MSP responsible for the service. A service agreement can establish what monitoring and response services were purchased, while operational reports can show that the service is actually functioning. Monthly security reports, alert and incident records, response reports, escalation procedures, and documentation describing monitoring hours can all support the answer.

The name of the MDR provider should also be verified rather than inferred from the EDR product. In some environments, the company providing the monitoring service is different from the company that makes the endpoint technology.

What the Executive Should Take Away from Endpoint Security

The easiest way for a business leader to understand this section is to focus on coverage and response. The company needs to know which devices it has, whether the required protections are actually installed on those devices, and what happens when one of those protections detects a potential attack.

This is an area where a simple conversation with IT or an outside provider can expose an important gap. Leadership may have been told that the company has EDR and reasonably believe that means someone is watching it around the clock. That may not be the case. The organization may have excellent detection technology generating alerts without anyone being responsible for responding to them outside normal working hours.

When supporting the insurance application, evidence should therefore show more than the name of the products the company purchased. It should establish that the expected devices are protected and, when the organization represents that it has managed detection and response, that someone is actually responsible for monitoring and responding when the technology says something is wrong.

Network Security

Network security questions help the insurer understand how the business controls the connections between its systems and the outside world, as well as what can communicate inside the company's environment. This is important because once an attacker gains access to one device or system, the amount of access available from that point can have a major effect on how serious the incident becomes.

For business leaders, network security can sound more complicated than it needs to be. The insurer is not expecting the person completing the application to understand firewall rules or network architecture. The practical question is whether the company has reasonable boundaries around its technology and whether someone is actively responsible for maintaining them.

Do you have a firewall in place for your network?

A firewall controls what network traffic is allowed into and out of the company's environment. It is one of the basic protections between business systems and the internet. Modern firewalls can also provide additional capabilities that identify suspicious connections, restrict access to certain services, support remote access, and provide information about activity occurring across the network.

The wording of this question matters because the application asks whether a firewall is in place for all networks. A company with one office may have a fairly simple answer. A company with several offices, warehouses, retail locations, production facilities, or other operating locations may have several networks that need to be considered. There may also be cloud environments or other technology that does not sit behind the firewall people normally think about at the corporate office.

The insurer is trying to determine whether the company's networks have an appropriate boundary between internal systems and external networks. Simply knowing that there is a firewall appliance in the server room does not completely answer that question. Someone needs to confirm that the organization's relevant networks are actually protected and that the firewall is being maintained.

The follow-up question asking for the firewall provider helps establish what technology is being used. The product itself is only part of the answer. A well-known firewall that has not been updated or whose configuration has not been reviewed for years can create a very different level of risk than leadership may assume.

IT, the network administrator, security team, or MSP will normally provide the information needed to answer this question. If different locations are managed by different providers, more than one person may need to be involved.

Evidence could include network diagrams showing where firewalls are deployed, firewall inventory records, configuration reports, screenshots from the management platform, maintenance or support records, and documentation from the MSP responsible for managing the devices. If the organization has several locations, the evidence should make it possible to confirm that the locations represented as protected actually have firewalls in place.

Is the guest network separate from the corporate or production network?

Guest Wi-Fi seems like a relatively minor issue, but this question is really about separation. Visitors, contractors, employees' personal devices, and other unmanaged devices may need internet access while they are at a company location. They generally do not need access to the same network used by company computers, servers, operational systems, or other important technology.

The insurer wants to know whether those guest devices are isolated from the business environment. If someone connects an infected personal laptop to guest Wi-Fi, that device should not suddenly have a path to corporate systems. The same concern applies to devices the company does not manage and therefore cannot reasonably know are secure.

For organizations with several locations, the word "all" again becomes important. The headquarters may have a properly separated guest network while a smaller office or branch uses a single wireless network for everyone. If the application says that the guest network is separate from all corporate networks, the organization should verify that this is true everywhere covered by the answer.

It is also possible that the business does not provide guest network access at all. That should be explained rather than forcing the situation into an answer that does not accurately describe the environment.

IT, the network administrator, wireless network administrator, or MSP should be able to verify how guest access is configured. Evidence could include network diagrams, wireless network configurations, firewall rules, VLAN configurations, or documentation showing that guest users receive internet access without being able to connect to corporate resources.

The executive reviewing the application does not need to understand what a VLAN is or how the firewall rules accomplish the separation. Someone responsible for the network should simply be able to demonstrate that a device

connected to the guest network cannot access the company's internal environment.

What the Executive Should Take Away from Network Security

These questions are really asking whether the business understands the boundaries around its technology. Leadership should know that the company's networks are protected from unnecessary internet access and that devices that do not belong on the corporate network are kept separate from it.

The evidence should also reflect the entire business rather than the location everyone happens to know best. This becomes particularly important when companies grow. New offices are opened, businesses are acquired, temporary locations become permanent, and different technology providers may be used over time. It is surprisingly easy for leadership to believe that every location follows the same security model when nobody has actually verified it.

For the insurance application, the person approving the answers should be able to confirm that someone responsible for the network has reviewed them and can support them with current information. The goal is not for an executive to become a network engineer. It is to make sure that a broad statement such as "a firewall is in place for all networks" is based on something more reliable than assumption.

Patch Management

Patch management is one of those areas that sounds like routine IT maintenance but has a direct connection to cyber insurance losses. Software vulnerabilities are discovered constantly, and once a serious vulnerability becomes public, attackers may begin looking for organizations that have not fixed it. The insurer wants to understand whether the business has a reliable way to identify those vulnerabilities and address them before they become an easy way into the organization.

For business leaders, the important distinction is between having automatic updates turned on somewhere and having an actual process for managing vulnerabilities across the business. Computers may receive Windows updates automatically while servers, firewalls, applications, cloud systems, mobile devices, or other technology follow completely different processes. The insurer is trying to understand how consistently the organization manages that exposure and how quickly it reacts when something particularly serious is discovered.

Do you have a process to patch or address critical vulnerabilities?

The insurer is asking whether the organization has a repeatable way to deal with serious security weaknesses after they are discovered. Patching is one way to address a vulnerability, but it is not always the only option. Sometimes a vendor has not released a patch yet, a system cannot immediately be updated for operational reasons, or the vulnerable feature can be disabled or isolated until a permanent fix is available. What matters is that someone identifies the problem, determines what systems are affected, and decides about what needs to happen.

The follow-up question about whether critical vulnerabilities are addressed within 7, 30, or 90 days is intended to determine how long serious weaknesses are likely to remain exposed. The answer should reflect what the company actually achieves, not simply what an internal policy says should happen. If company policy requires critical patches within seven days but reports routinely show that systems remain unpatched for several weeks, answering seven days would create a representation that the evidence does not support.

The word critical also deserves some attention. Not every missing software update presents the same risk. A vulnerability that can be exploited remotely on an internet-facing system may require much faster attention than a vulnerability on an isolated internal device. The organization's process should provide some way of determining which vulnerabilities require immediate attention rather than treating every update exactly the same.

The application then asks where the patching process applies. This is important because many organizations have good patch management for employee laptops and desktops but a less consistent process for servers, applications, mobile devices, network equipment, or other systems. Selecting that patches are rolled out across all endpoints or servers should be based on actual coverage.

The question about secure or hardened baseline configurations looks slightly beyond patching. The insurer wants to know whether systems are also being maintained against an established configuration rather than simply receiving software updates. A system can be completely patched and still be poorly secured because unnecessary services, accounts, or insecure settings remain enabled.

IT, the system administrator, security team, vulnerability management team, or MSP will normally provide this information. Evidence can include patch compliance reports, vulnerability scan results, endpoint management reports, server patching reports, remediation tickets, configuration management reports, and documentation showing the organization's expected patching timeframes.

One of the strongest forms of evidence is a report that shows both the systems being managed and their current patch status. A written policy saying that critical vulnerabilities are patched within seven days establishes the expectation. A report showing that the organization actually meets that expectation establishes what is happening in practice.

Do you have a Managed Service Provider (MSP) or Managed Security Services Provider (MSSP)?

This question appears under patch management because many organizations rely on an outside provider to perform some or all of their technology maintenance. For smaller businesses in particular, an MSP may be responsible for managing employee computers, deploying patches, maintaining servers, administering Microsoft 365, managing firewalls, or performing other day-to-day technology functions.

An MSSP generally has more of a security focus and may provide vulnerability management, security monitoring, incident response support, managed endpoint security, or other security services. The distinction is not always clean because many providers now offer both IT and security services.

For the insurer, the important issue is not simply whether the company has hired an MSP or MSSP. It is what that provider is actually responsible for. Business leaders sometimes assume that because an outside company "handles IT,"

patching and vulnerability management must be included. That may not be what the contract says, and it may not be what the provider is doing.

If the business relies on an outside provider for patch management, someone should understand which systems the provider manages, how frequently patches are deployed, what happens when a critical vulnerability is discovered, and who is responsible for systems outside the provider's scope. This is especially important when several providers are involved. One company may manage employee laptops while another manages a business application, and someone internally remains responsible for servers or network equipment.

The name of the MSP or MSSP should be easy to establish, but the supporting evidence should go further. The service agreement, statement of work, or contract can show what the provider is responsible for. Patch compliance reports and service reports can then demonstrate that the work is actually occurring. If the provider is expected to address critical vulnerabilities within a specific timeframe, the agreement or operating procedures should make that responsibility reasonably clear.

Finance, procurement, or the business owner may know who the provider is and hold the contract, while IT or the provider itself will usually have the operational evidence. This is a good example of why insurance applications often require information from more than one part of the company.

What the Executive Should Take Away from Patch Management

The executive reviewing this section does not need to know which Microsoft update was installed last Tuesday. The important thing is knowing whether the organization has a process that can answer some basic business questions when a serious vulnerability appears. Someone should know whether the company is affected, which systems need attention, who is responsible for fixing them, and whether the problem was actually resolved.

This section can also expose a common gap between outsourcing responsibility and assuming responsibility has been outsourced. Hiring an MSP does not automatically mean every system is patched or every vulnerability is being managed. The contract and operational reports should show what the provider actually does, while the business should understand who owns anything that falls outside that scope.

When the insurance application asks whether critical vulnerabilities are addressed within a certain number of days, that answer should ultimately be supported by evidence of actual performance. That is considerably more useful to the business,

and safer for the person approving the application, than relying on what everyone assumes the patching process is supposed to do.

Data Backup

Data backup is where the cyber insurance application begins moving beyond whether the company can prevent or detect an attack and starts looking at what happens when prevention fails. Insurers care about backups because the length and cost of a cyber event can change dramatically depending on whether the business can recover its systems and information. A ransomware event that can be restored from reliable backups is a very different insurance loss from one where the backups were also encrypted, damaged, or simply did not work.

For business leaders, this section deserves more attention than asking IT whether the company has backups. Most organizations do. The more useful questions are what is being backed up, where those backups are stored, whether an attacker could reach them, and whether the business knows they can actually be restored. A successful backup report tells you that data was copied somewhere. It does not necessarily tell you that the business can recover from it.

Do you have a backup process for your data?

The insurer first wants to establish whether important company data and systems are regularly backed up and whether another usable copy exists if the primary environment becomes unavailable. The distinction in the application between data simply being backed up and being stored securely at an alternate site is important. If the only backup is located alongside the original systems and accessible through the same environment, a ransomware attack or other major event could potentially affect both.

An alternate site does not necessarily mean another physical building. Many organizations use cloud-based backup services or geographically separate data centers. What matters is whether the backup environment provides meaningful separation from the systems it is intended to recover.

The follow-up questions reveal more about what the insurer is really trying to determine. Asking whether sensitive data and critical business systems are backed up at least weekly addresses how much information could be lost. For many businesses, weekly backups may still leave an unacceptable amount of missing information, so the organization's actual backup frequency should be understood rather than simply checking that it meets the minimum presented by the application.

The question about testing backups is particularly important. A backup can complete successfully every night and still fail when someone tries to restore from it. Files may be corrupted, an application may require additional components that

were never backed up, or the organization may discover that the recovery process takes far longer than expected. Periodic restoration testing provides evidence that the backup can actually be used.

The insurer also asks whether backups are disconnected or inaccessible from the primary network because ransomware operators increasingly try to damage backups before attacking production systems. If an attacker compromises an administrator account and that same account can delete or encrypt the backups, the existence of those backups may provide much less protection than leadership believes.

The question about having an authentication mechanism outside LDAP or Active Directory addresses a similar problem. If the organization's primary identity environment is compromised, the insurer wants to know whether the company can still reach its recovery environment. A backup that depends entirely on the same compromised credentials may become unavailable exactly when it is needed.

IT, the backup administrator, infrastructure team, cloud administrator, MSP, or backup service provider will normally have the information needed to answer these questions. Different systems may also have different owners. The company should make sure the answer covers the critical systems that actually operate the business rather than assuming that one backup product protects everything.

Evidence can include backup job reports, configuration reports showing backup frequency and retention, documentation showing where backups are stored, access configurations, and records from restoration tests. Reports showing successful backup jobs are useful, but evidence of successful restoration is considerably more meaningful when the application asks whether backups are periodically tested.

Do you have a plan to restore business functionality and data in case of a disaster?

This question goes beyond backups. Having the data needed to rebuild the business is important, but the organization also needs to know how it would actually put everything back together.

A disaster recovery plan should explain how important technology is restored after a significant disruption. That may involve recovering servers, cloud environments, applications, databases, identity systems, network services, and the information employees need to work. The order matters because systems often depend on one another. Restoring a business application may accomplish very

little if the identity service, database, network connection, or other technology it depends on is still unavailable.

The distinction between having a written plan and having a tested plan is therefore significant. A document may describe what everyone believes should happen. Testing provides evidence about whether it actually can happen. A recovery exercise may discover missing passwords, outdated instructions, unavailable personnel, forgotten system dependencies, or recovery times that are considerably longer than leadership expected.

For an executive or financial leader, this question has a direct connection to business interruption. If a major cyber event shut down the company's critical technology tonight, how long would it take before the business could operate again? That period can determine how much revenue is lost, how long employees remain unable to work, and how large the resulting insurance claim becomes.

IT, business continuity personnel, disaster recovery teams, application owners, MSPs, and cloud providers may all contribute to the answer. Leadership and business operations should also be involved because IT cannot determine by itself which business functions must return first or how long the company can reasonably operate without them.

Evidence should include the disaster recovery plan itself along with records showing when it was reviewed and tested. Test reports should document what was restored, whether the recovery worked, what problems were discovered, and whether those problems were subsequently addressed. If the organization has actually performed system restoration exercises, those records provide much stronger evidence than a plan that has never been used.

What the Executive Should Take Away from Data Backup

This section should not end with someone telling leadership, "Yes, we have backups." The executive reviewing the insurance application should have enough information to know whether the company's critical systems and data are covered, whether the backups are reasonably protected from the same event affecting the primary environment, and whether anyone has demonstrated that they can actually be restored.

There is also an important business question hiding behind the technical language. The company may be able to restore its data and still be unable to operate for several days. That is why the disaster recovery question matters just as much as the backup question. Recovery is ultimately about restoring the business, not simply restoring files.

For the insurance application, the evidence should support what the company says it can do today. For leadership, it should provide something more useful: reasonable confidence that if the systems the business depends on disappeared tomorrow, the organization has a tested way to get them back.

Incident Response

Incident response is where the insurance application starts asking what the organization will actually do when a cyber event occurs. Up to this point, many of the questions have focused on reducing the chance of an incident or limiting an attack once it begins. Incident response assumes that something has already happened and looks at whether the business can make good decisions quickly enough to keep the situation from becoming worse.

This matters to an insurer because the first several hours of a cyber event can have a significant effect on the eventual cost of the claim. A company that knows who is responsible, how to reach its insurance carrier, when to involve legal counsel, and how to bring in outside specialists is in a much different position from a company trying to figure those things out while systems are already unavailable. For executives, this is also where cyber risk becomes much less of an IT issue. A serious incident can quickly involve leadership, finance, legal, operations, communications, insurance, customers, and outside authorities.

Do you have an incident response plan?

An incident response plan describes how the organization will manage a cyber incident once it has been identified. It should establish who becomes involved, who has authority to make decisions, how the situation is escalated, and how the organization coordinates the technical and business response.

The insurer is interested in whether the plan is documented because an informal understanding that "IT will handle it" tends to break down during a serious event. The people normally responsible for technology may themselves be dealing with unavailable systems, compromised accounts, or incomplete information. Leadership may suddenly need to make decisions about shutting down systems, notifying customers, involving law enforcement, engaging forensic specialists, or activating the cyber insurance policy.

The distinction between a documented plan and a documented and tested plan is important. A company can have an excellent-looking incident response document that has never been used or exercised. Testing the plan gives the organization an opportunity to discover whether the people listed in it still work for the company, whether contact information is current, whether everyone understands their role, and whether important decisions can actually be made under realistic conditions.

Testing does not necessarily require deliberately disrupting the company's systems. A tabletop exercise can walk leadership and the response team through a realistic scenario and require them to explain what they would do as the situation

develops. The value comes from identifying confusion and missing information before those problems appear during a real incident.

IT and security will normally own part of the plan, but the answer should not come from them alone. Legal, finance, operations, communications, senior leadership, and the company's insurance broker or risk function may all have responsibilities during an incident. Outside providers may also be involved, particularly when the company relies on an MSP, MDR provider, incident response firm, or outside legal counsel.

Evidence can include the current incident response plan, records showing when it was reviewed, tabletop exercise documentation, test results, attendance records, and records showing that problems identified during testing were addressed. If the organization answers that the plan is tested, there should be some reasonable evidence showing when that testing occurred and what was actually exercised.

Are incidents logged and tracked?

This question asks whether the organization maintains a record of security incidents and follows them through resolution. The insurer wants to understand whether incidents are being managed through a defined process or handled informally through emails, phone calls, and conversations that may leave little record of what actually happened.

The word "all" matters here. Selecting that all incidents are logged and tracked until resolution is a broader statement than saying the company uses a ticketing system. The organization should understand what it considers a security incident, how those incidents enter the process, and whether they remain tracked until someone determines that the issue has been resolved.

Good incident records can also become important during an insurance claim. They may help establish when the organization first detected suspicious activity, what actions were taken, who was involved, and how the incident developed. That history can be difficult to reconstruct after the fact if nobody maintained a reliable record while the event was occurring.

The security team, IT department, SOC, MDR provider, or MSP will normally maintain these records. Depending on the organization, incidents might be tracked through a service management platform, security operations platform, case management system, or another formal process.

Evidence can include incident tickets, case records, security incident reports, SOC reports, MDR reports, and documentation describing the company's incident tracking process. The organization does not need to provide an insurer with

sensitive details about every incident simply to prove that a process exists. A report showing incident volumes, status, dates, and closure information may provide appropriate evidence without unnecessarily exposing confidential information.

Is there a process for reporting incidents to appropriate authorities?

This question reaches beyond the technical response and into the company's legal, regulatory, contractual, and insurance responsibilities. Depending on what happened, an organization may need to notify its insurer, customers, regulators, affected individuals, law enforcement, business partners, or other parties.

The insurer wants to know whether the company has thought about those obligations before an incident occurs. Determining who must be notified can become complicated very quickly when personal information has been exposed, several jurisdictions are involved, or contracts contain specific notification requirements. That is not something leadership should be trying to understand for the first time in the middle of a major breach.

The phrase "appropriate authorities" should also be interpreted carefully. Not every cyber incident needs to be reported to a regulator or law enforcement agency. The organization should have a process for determining whether notification is required and who is authorized to make that decision. Legal counsel will often play an important role in making that determination.

The cyber insurance policy itself should be part of this process. Policies typically contain requirements describing how and when the insurer should be notified of a potential claim. They may also provide access to approved legal counsel, forensic investigators, ransomware negotiators, notification providers, and other specialists. Engaging outside firms without understanding the policy requirements can create unnecessary complications over whether those expenses are covered.

Legal or compliance personnel will normally provide much of the information needed for this answer, with involvement from leadership, risk management, finance, IT, and the insurance broker. Smaller businesses that do not have internal legal counsel should understand who they will contact for guidance when an incident occurs rather than waiting until they are already dealing with one.

Evidence can include the notification section of the incident response plan, regulatory notification procedures, escalation documentation, current cyber insurance contact information, broker and carrier notification procedures, and contact information for outside counsel or other response partners. If the

organization has previously conducted a tabletop exercise, the exercise should ideally include the question of when the insurer and other outside parties would be contacted.

What the Executive Should Take Away from Incident Response

This section is ultimately asking whether the business is prepared to manage the decisions that follow a cyber event. Having security technology does not answer that question. When an actual incident occurs, someone has to determine what happened, decide what needs to be done, involve the right people, document the response, and understand whether anyone outside the company needs to be notified.

Executives should pay particular attention to the statement that the incident response plan has been tested. A plan sitting in a shared drive is evidence that someone wrote a plan. It does not establish that the organization can execute it. Leadership should know when the last exercise occurred, who participated, what was learned, and whether the issues discovered during that exercise were addressed.

There is also a direct insurance connection that should not get lost in the technical discussion. The cyber insurance carrier should already be part of the company's response thinking before a claim happens. Leadership should know where the policy is located, who can notify the carrier, who the broker contact is, and what the policy requires before outside response services are engaged. When the business is already dealing with ransomware or a major outage is the wrong time to start looking for the policy.

Personnel Security

Personnel security looks at what happens before someone is given access to the business and what happens when that access should end. These questions may seem more like HR questions than cybersecurity questions, but from an insurance perspective they deal with a very practical risk: employees and contractors often have legitimate access to systems, data, email, financial information, and customer records. The organization needs a reasonable process for managing that access throughout the employment relationship.

Are background checks conducted for employees?

This question asks whether the organization performs background screening before giving employees access to company systems, facilities, money, or sensitive information. The insurer is looking for evidence that the company has some process for evaluating personnel risk rather than automatically providing access based solely on hiring.

The answer should reflect what the organization actually does. There is an important difference between conducting background checks for all employees and conducting them only for certain positions. If the company screens executives, finance personnel, IT administrators, or employees handling sensitive information but not every employee, it should select the second answer rather than treating the broader answer as the "better" insurance response.

HR will normally own this process, although legal or compliance may establish requirements for particular positions. Smaller businesses may use an outside screening company as part of their hiring process.

Evidence does not need to include employees' actual background reports, and I would generally avoid uploading those because they contain sensitive personal information. Better evidence would be the company's hiring or screening policy, documentation showing that background screening is required, or a report from the screening provider demonstrating that the process is being performed without exposing individual results.

One caution with the wording in the application: the Fair Credit Reporting Act regulates how consumer reports may be obtained and used for employment purposes. It should not be read as a blanket requirement that every organization conduct background checks on every employee.

Is there a process for handling employee terminations?

From a cybersecurity and insurance perspective, this is the more important of the two questions. When someone leaves the company, their access should leave with them.

A formal termination or offboarding process should connect HR with whoever controls technology access. That may be an internal IT department, MSP, identity administrator, security team, or some combination of them. When HR determines that someone's employment is ending, there should be a defined way to tell the people responsible for disabling that person's access.

The process should cover more than simply turning off the employee's primary email account. Depending on the person's role, access may exist in Microsoft 365 or Google Workspace, VPNs, cloud services, financial systems, CRM platforms, shared accounts, remote-access tools, privileged administrator accounts, physical access systems, and third-party applications. Company laptops, phones, security tokens, badges, keys, and other equipment may also need to be recovered.

Timing matters. For a normal departure, access should be removed according to the established offboarding process. For an involuntary termination or a situation involving elevated risk, HR and IT may need to coordinate access removal with the termination itself so that the employee does not retain system access after being notified.

This is also where many organizations discover a weakness that isn't obvious from the application. Removing an Active Directory or Microsoft 365 account does not necessarily remove every account that person accumulated while working for the company. SaaS applications, shared credentials, vendor portals, cloud environments, and locally managed systems can remain accessible unless the organization maintains a reasonable understanding of who has access to what.

Evidence can include the employee offboarding procedure, termination checklist, HR-to-IT workflow, service tickets showing completed access removal, equipment-return records, or reports showing disabled accounts. As with background checks, evidence should demonstrate the process without unnecessarily providing sensitive employee information.

What the Executive Should Take Away from Personnel Security

These two questions are really asking whether the organization manages the human side of access. Background screening addresses risk before access is granted. The termination process addresses what happens when that access is no longer appropriate.

The second question deserves particular executive attention because former-employee access is usually preventable. Leadership should be able to ask a simple question after someone leaves: How do we know they can't still get into anything?

A good answer isn't "HR told IT." The business should be able to show that the notification happened, the relevant access was identified, accounts were disabled or removed, company assets were recovered, and the process was completed. That turns an HR procedure into an actual security control.

Can We Insure This?

Working through these questions makes something pretty clear. A cyber insurance application isn't really asking whether you have a collection of security products. It's trying to determine how much confidence an insurer should have in the way your business actually operates.

Access control, employee awareness, email security, endpoint protection, network security, patching, backups, incident response, and personnel security all contribute to that picture. The individual questions can look simple. Yes or no. Check a box. Upload some evidence. But taken together, they tell a much bigger story about how prepared the business is when something goes wrong.

That's one of the central ideas behind Can We Insure This? 2nd Edition. Cyber insurance has moved well beyond filling out an application once a year and hoping for a reasonable premium. Underwriters increasingly want to understand whether the controls you're claiming actually exist, whether they work consistently, and whether the business can recover when those controls fail.

There is also an important distinction between being insurable and being survivable. You can have MFA, EDR, firewalls, backups, an incident response plan, and a cyber insurance policy and still discover during an incident that the business can't operate. A backup that has never been restored isn't much comfort during ransomware. An incident response plan nobody has exercised may fall apart under pressure. An insurance policy can provide financial protection, but it can't operate the company while systems are unavailable.

This is why I don't think executives should treat these questions as something to hand to IT and forget about. They provide a useful window into the business itself. If you can't confidently answer a question, can't prove the answer, or discover that the answer depends entirely on one employee or vendor, you've probably found something worth understanding regardless of what the insurance application requires.

That brings us back to the question behind the book: Can we insure this? The answer increasingly depends on another question that matters even more: If something happens, can the business survive it?

Can We Insure This? 2nd Edition: A Business Leader's Guide to Cyber Risk, AI, Insurance, and Business Survivability explores that connection in much greater detail, helping business leaders understand what cyber insurance actually covers, what insurers are looking for, where coverage gaps can emerge, and why insurance should be part of a broader approach to business survivability.