Assume breach: Managing a dirty network

Contents

Foreword

In December 2020, it became apparent that SolarWinds, a major US information technology firm, had been the subject of a cyber attack that spread to its clients and went undetected for months.

The SolarWinds incident has had a huge impact across the entire technology ecosystem and it continues to unfold. It's unlikely that we'll know the full extent of its effects for many more months, and the ramifications of the incident will continue for even longer.

It's had one immediate effect though: the whole security community is now questioning some of the fundamental practices and assumptions that it makes around how to implement a successful security environment. Attackers have been more successful, on a wider scale and across a longer timeframe, than many within the security industry thought possible, given the protections many organisations already had in place.

This raises several questions around how organisations respond. How do you secure your organisation and data on a breached or 'dirty' network?

How do you put strategies and technologies in place to cope with the fact that you may have already been successfully attacked, and just not know it yet?

More widely, the SolarWinds attack is forcing a rethink of how to assess and manage supplier risk, and the trusted access those suppliers have to critical systems and information. All security professionals know that you're only ever as strong as the weakest link in your defences. Do you now need to evolve and strengthen 'Zero Trust'? Or must you simply assume a breach will happen somewhere amongst your suppliers eventually, and put in place more proactive and reactive defences against this?

This paper shares expert insight we've gained from across our broad client base and our own global infrastructure, as well as the guidance we're following to secure ourselves and our customers. This paper also looks at what you should consider – in light of this incident – when planning your future security strategy.

Kevin Brown,
Managing Director, BT Security

Introduction

It's increasingly difficult to gain full control and transparency over an IT environment. Most organisations don't grow organically over time; they grow through a series of mergers, acquisition and divestments which each play a part in changing their IT landscape. As their infrastructure evolves, it becomes a mixture of new, established and legacy systems from a range of different suppliers.

Despite this, users expect IT to be frictionless, leading many organisations to become increasingly borderless. This makes it far harder to clearly define the lines between applications, user groups and even stakeholders.

In this complex, blurred environment, finding 'bad' - or even sub-optimal – elements can be challenging, because there are lots of places for 'evil' to hide.

Focus on the basics, not the flashy

It's easy to be distracted by the latest buzz in the industry, but often the most effective solutions lie in following established best practices. That's not to say implementation will be straightforward, but using simple controls and starting with the most abused capabilities will often yield the best results for protecting your business.

Act differently, assume the worst

You need to 'assume breach' – expecting a sophisticated attacker will find a way into your estate and manage to stay hidden for some time.

Assuming the attacker is already inside changes your goal to making it harder for them to move about - and easier for you to detect. Give them little room to cling on, and plenty of obstacles to trip them up. The SolarWinds incident has shown how important taking this security stance is.

The following six recommendations explore the policies and solutions we're following to achieve 'assume breach', and our recommendations on how other organisations can adopt them too.

Recommendation one: Know the personas on your estate (identity)

As IT has proliferated within our business and personal lives, so too have the number of accounts and associated privileges, with many applications requiring specific and standalone identities.

What's more, the type of access and permissions that users require are evolving more dynamically than ever before, and the task of tightly controlling 'least privilege' access can be difficult to keep on top of.

As we work towards greater automation, orchestration, and ultimately federation, Application Programming Interface (API) access and machine-to-machine accounts are adding to the problem.

The pattern of an attack

To an attacker, overly permissive accounts or being able to use a compromised device to move laterally around an organisation is very appealing.

A cyber criminal will follow a set path of attack, often referred to as the Cyber Kill Chain (Lockheed Martin®). Once inside an organisation, an attacker will look to:

- protect their access by creating alternative covert entry paths to the organisation's assets
- elevate their privileges to make it easier to move and hide within the organisation
- move laterally to find what they're looking for.

How to respond

The complexity of managing and understanding personas and identities leaves many organisations blind to the activity of an attacker.

In this context, identity and access mechanisms that give you visibility and control of your estate are hugely valuable.

These mechanisms allow organisations to identify when a user's privileges are elevated, when new administration accounts are added (or modified), and when application or federation access is created.

This is especially important for identifying stores and directories used by multiple applications, as well as services that have wider than usual trust levels within an estate.

Additional advanced techniques can also help by reporting on the creation and use of application interfaces and federation, which often seek to extend privilege to machines and other organisational entities outside your visibility and control.

Identity is one of the areas of compromise frequently implicated in high profile and impactful breaches. So a firm understanding of the roles and users in your organisation, coupled with high confidence audit, reporting and alerting, is critically important. This can go a long way to helping detect malicious behaviour more rapidly and limits the impact by ejecting an attacker earlier in their campaign.

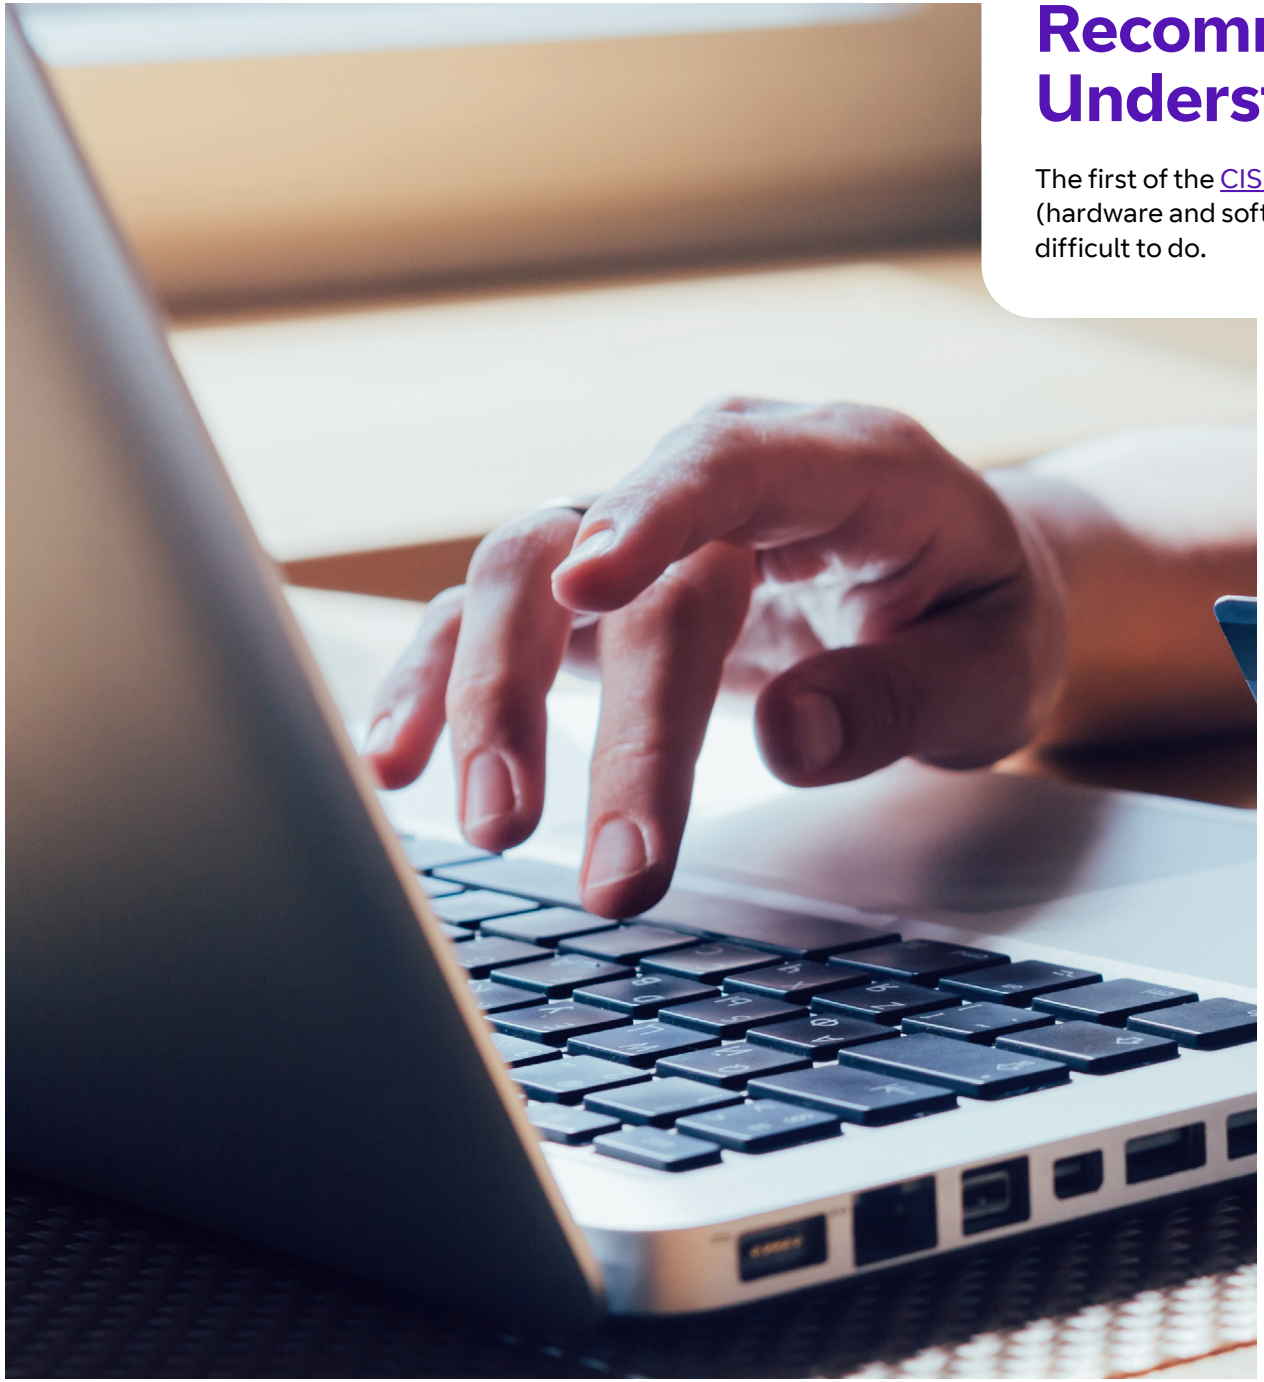

Recommendation two: Understand your assets

The first of the CIS 20 controls focuses on understanding your assets (hardware and software). Fundamental stuff, but still remarkably difficult to do.

Why is this such a problem?

It comes back to the old adage that if you don't know what you have, how can you protect it? The proliferation of IT and IoT devices, coupled with the hyper-connectedness of everything, means that it can be hard to fully identify your vulnerabilities and risks - meaning some remain undetected.

Not having a holistic view of your network and assets becomes a critical issue when a major security incident happens. In fact, "do we have one of those?" is the most common question when a security issue hits the news.

SolarWinds was particularly problematic for many organisations because it was the type of attack we call an 'enthusiast's tool'; something hidden under a sysadmin's desk quietly monitoring the network, and therefore not officially part of the known estate.

Unfortunately, many organisations don't know what should be on their network, let alone what is actually out there.

Keeping track of the reality of your assets is an ongoing task and there are a lot of discovery tools and several asset management systems available to help. It's never been more important to do it, to cope with all the unknown, unmanaged assets that are easy prey for attackers.

The asset circle of life

Understanding what and where your assets are is only one part of the problem. You also need to rigorously assess your asset life cycle strategy.

Often, when a serious vulnerability is identified, it can be very difficult to track down the assets affected, and hard to understand their patching or version status so that you can update and defend the device.

Breaking the asset lifecycle into three logical steps helps to optimise each step and makes it easier to carry out patching, should the worst happen.

1. **Know what you have**
   Creating a high-integrity asset or configuration database of what is connected to your infrastructure helps to rapidly assess the scope of vulnerabilities. Ideally, you'll update this whenever new devices are detected. A strong asset approach can also help in managing licenses and support costs, as well as in identifying end-of-life or legacy IT that may be hiding in your network.
2. **Know what is vulnerable**
   Regular scanning of your asset database to identify the versions of software in use and any potential vulnerabilities can help you frame the IT risk your organisation faces, and help to prioritise your remediation strategy.
3. **Resolve the risks**
   A solid approach to patching which makes the most of your asset knowledge will mean you can close vulnerabilities faster, while minimising the impact to your business operations.

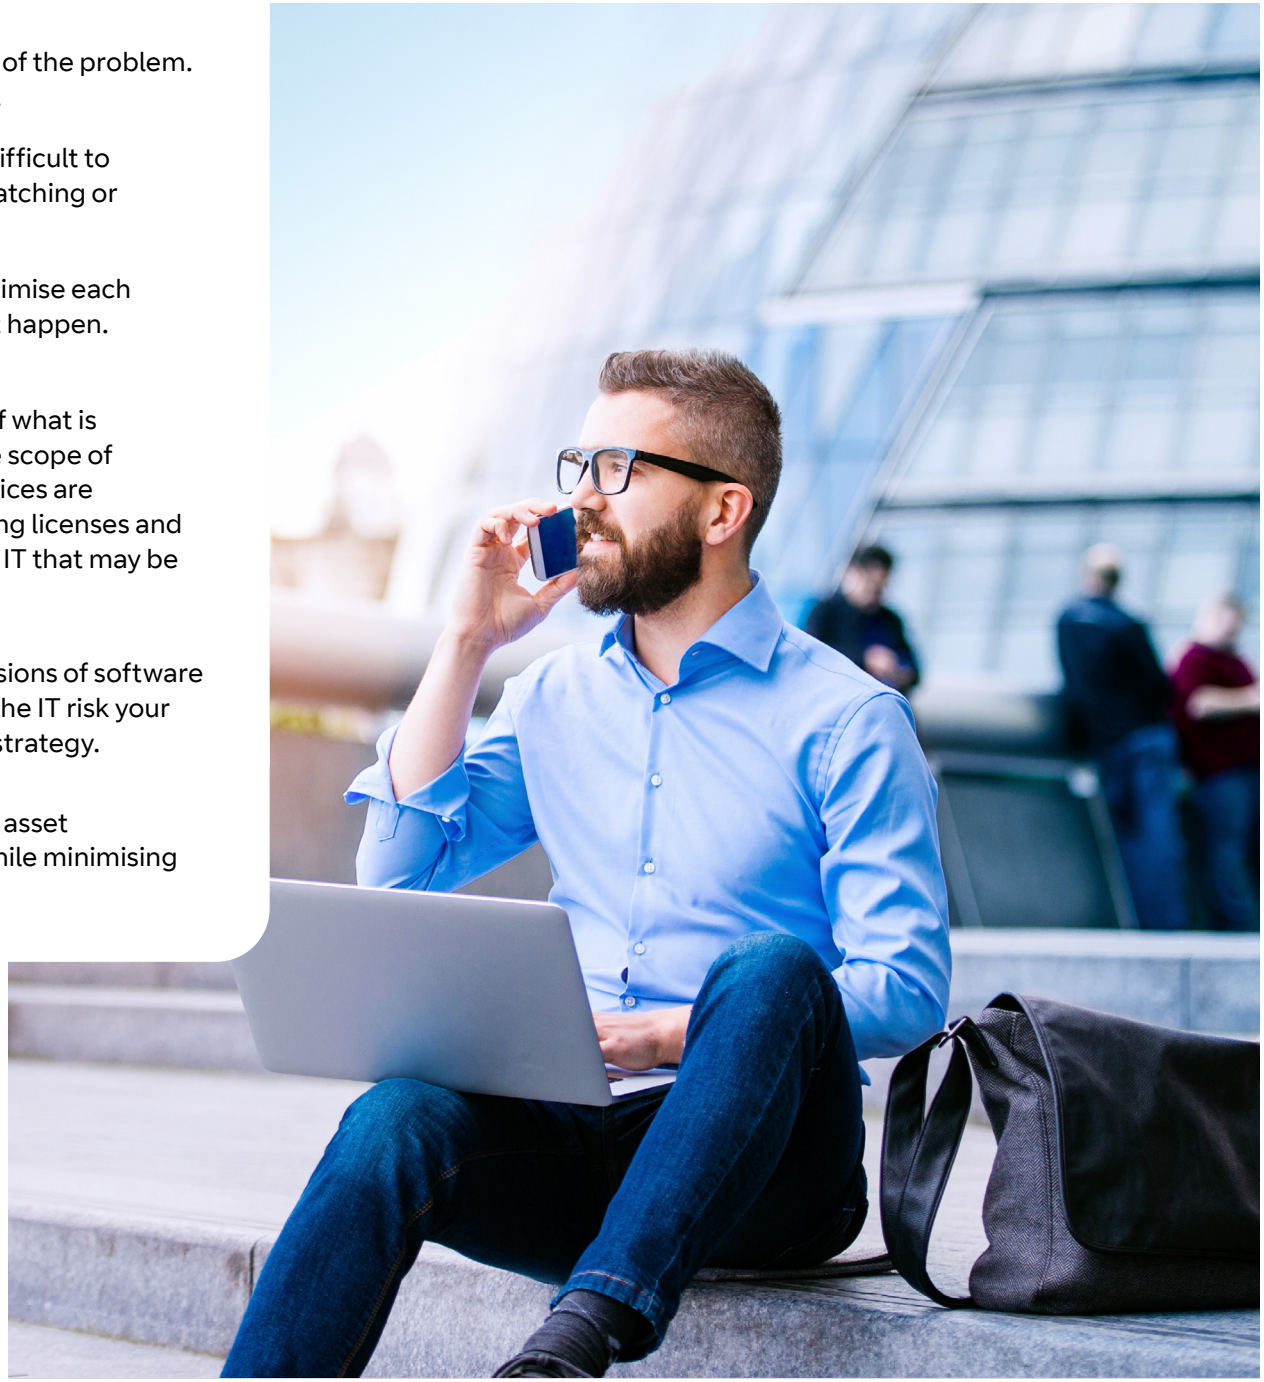

The asset lifecycle is perhaps one of the most difficult aspects of successfully managing IT infrastructure, and it gets more difficult as you move to the cloud, and more corporate assets fall outside your traditional network perimeter.

However, high profile incidents like SolarWinds (2020) and WannaCry (2017) demonstrate what's at stake if you fail to identify affected versions. They can delay the remediation and patching process, even after fixes are made available, worsening the risks and impacts.

Recommendation three: Prioritise modern endpoint (EDR) tooling

Challenging your security assumptions

Cyber attacks continue to get more sophisticated and difficult to defend against. Incidents such as SolarWinds create a new paradigm that challenges the basis of what we thought we knew and our understanding of who we trust. This has sent a seismic shock through the security industry, causing security teams to review what this new reality means to them.

However, using events like SolarWinds as opportunities to test the viability of your existing protection strategies is no bad thing. In fact, using real-world examples to stress test how successfully you would defend against similar incidents is an excellent way to improve your security posture.

These exercises can uncover where you've added more and more countermeasures into your security architecture, with decreasing returns. In fact, this approach is unsustainable because it increases complexity and ramps up the workload of already stretched security teams. New tools and systems add to management responsibilities, increase data volumes, and lead to more alerts that need investigating. Many organisations struggle to keep up. But there is another way.

How EDR can tackle this problem

Endpoint Detection and Response (EDR) solutions tackle this issue by bringing together next-generation antivirus with threat hunting and threat intelligence on the endpoint device, constantly analysing events to identify malicious behaviour.

Although an EDR solution gives excellent visibility of adversary behaviour as it occurs, organisations often need prior understanding of this behaviour to detect and prevent it effectively. When this information isn't available, and prevention and detection fail silently (as was seen with the SolarWinds incident), many EDR solutions monitor and record the chain of execution of activities occurring on the endpoint. This tool can provide visibility of what's happening now, and retrospectively, meaning that when details of an attack are made available, SOC teams can look back and verify where the attack happened.

Recommendation four: Make it difficult to move between zones and workloads

It's time to go beyond flat

Organisations need to move beyond the traditional, flat logical network architecture that allows every device to communicate with whatever they need to. Instead, organisations must adopt a Zero Trust model that's secure by default, and only allows traffic to flow between applications that's been positively verified against policy. This will reduce the opportunity for malware or threat actors to move between network zones, servers or workloads, providing crucial protections during many cyber incidents.

Creating boundaries between different zones of your network, using network segmentation and application micro-segmentation can make it more difficult for an attacker to move laterally around your infrastructure. It effectively creates 'fire door' control points that can rapidly restrict access if needed, and creates logical boundaries for establishing tripwires or inspection points to detect anomalies.

There are several quality network and micro-segmentation tools on the market that can help you achieve this. Introducing them as part of re-architecting your data centre topologies and network architectures will play a huge part in restricting movement around your organisation and in securing your data.

Cloud makes it easier, but brings challenges too

The cloud delivers huge benefits to organisational flexibility, scale and delivery, but it also adds to the security complexity you need to tackle.

The good news is that whether you're talking network or application security groups in Azure, or VPC firewall rules and network clusters in Google Cloud, each of the cloud providers has native tooling to segment workloads and compute tasks.

However, you've still got to set these things up. It's easy to ignore, thinking "we'll get to that later" – or to start

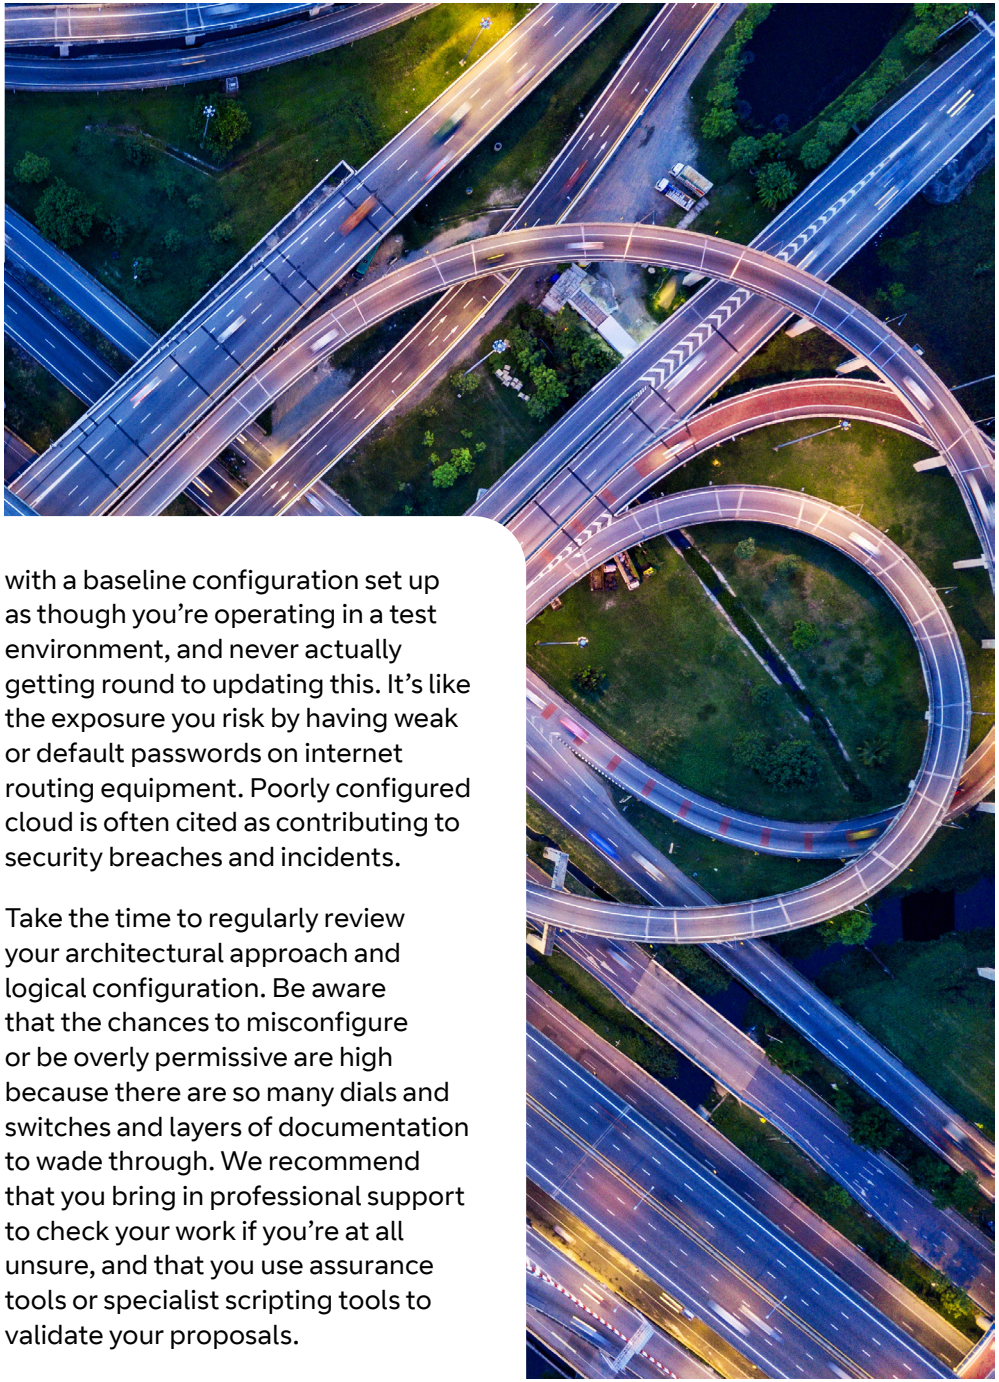

with a baseline configuration set up as though you're operating in a test environment, and never actually getting round to updating this. It's like the exposure you risk by having weak or default passwords on internet routing equipment. Poorly configured cloud is often cited as contributing to security breaches and incidents.

Take the time to regularly review your architectural approach and logical configuration. Be aware that the chances to misconfigure or be overly permissive are high because there are so many dials and switches and layers of documentation to wade through. We recommend that you bring in professional support to check your work if you're at all unsure, and that you use assurance tools or specialist scripting tools to validate your proposals.

Recommendation five: Take a systemic approach to detecting threats

Spotting what's out there

Organisations invest in threat detection capabilities such as SIEM (Security Information and Event Management) to make sure they can detect compromises within their estate quickly.

However, detecting threats effectively means deploying and tuning the analytical capability behind it properly. It's not a simple undertaking for any organisation.

During times of heightened awareness of cyber threats (which often take place when there are public compromises in the media), organisations ask themselves how much protection they are getting from their security investments. SIEM is particularly valuable at times like these because, if it's working effectively, the SOC team can demonstrate that the organisation hasn't been compromised.

Fine tuning your detection

To achieve this, the SOC team operating the SIEM needs to adopt a systematic approach.

First off, they need a good understanding of threat actor behaviour. SOC teams should work closely with their counterparts in threat intelligence to identify the behaviour of known actor groups and to map this knowledge to a common classification structure, such as the Mitre ATT&CK® framework.

Put simply, this means your teams can develop a better understanding of the activities they're trying to detect, making their jobs much easier (and more effective). Adding current detection capabilities onto this mapping of adversary behaviour shows the coverage they have and, importantly, where there are gaps.

The SOC team can then work with the business to increase this coverage through gathering data and the development of suitable analytics. Once they understand the threat detection coverage, other stakeholders can make key decisions about whether the risk associated with current gaps is acceptable, or if investment is needed to reduce the risk.

This mapping process lets the SOC team determine how effectively they can detect threats, and helps them prioritise any remediation activities in a targeted and strategic way.

Recommendation six: Be curious

Use the human firewall

Tools and processes are great, but people are better. Encouraging your end users to be curious about how they can support your organisation's security aims, and providing them with the skills they need to protect you, will bring significant benefits.

Unfortunately, those same people can be your biggest weakness. Users naturally want to do the right thing, and will look for ways to follow your processes and procedures. But if you make it too hard for them, or don't give them a mechanism to follow, then they'll do their own thing and unknowingly create risks.

Find ways to engage with users, and make sure you have the right tools in place to support them. Informed users are your best way to find things that are out of place. If you put clear pathways in place for reporting issues, backed up by a positive and supportive culture, you'll have extra eyes and ears helping you spot vulnerabilities and defend the organisation.

Give your analysts room to explore

The most inquisitive and engaged people in the organisation are the analysts you have defending your estate. So, make the most of them by managing their workloads and automating volume activity where possible, so they can focus on using their natural talents to maximum effect. Burdening them with repetitive or routine tasks might produce a steady flow of outputs, but it isn't the most effective use of their time or skills.

Consider automating or offloading such items to trusted providers, so your analyst can better spend their time being curious or searching things out. Pulling on loose threads takes time, but ultimately, it improves your security baseline and might just uncover the thing no-one was looking for.

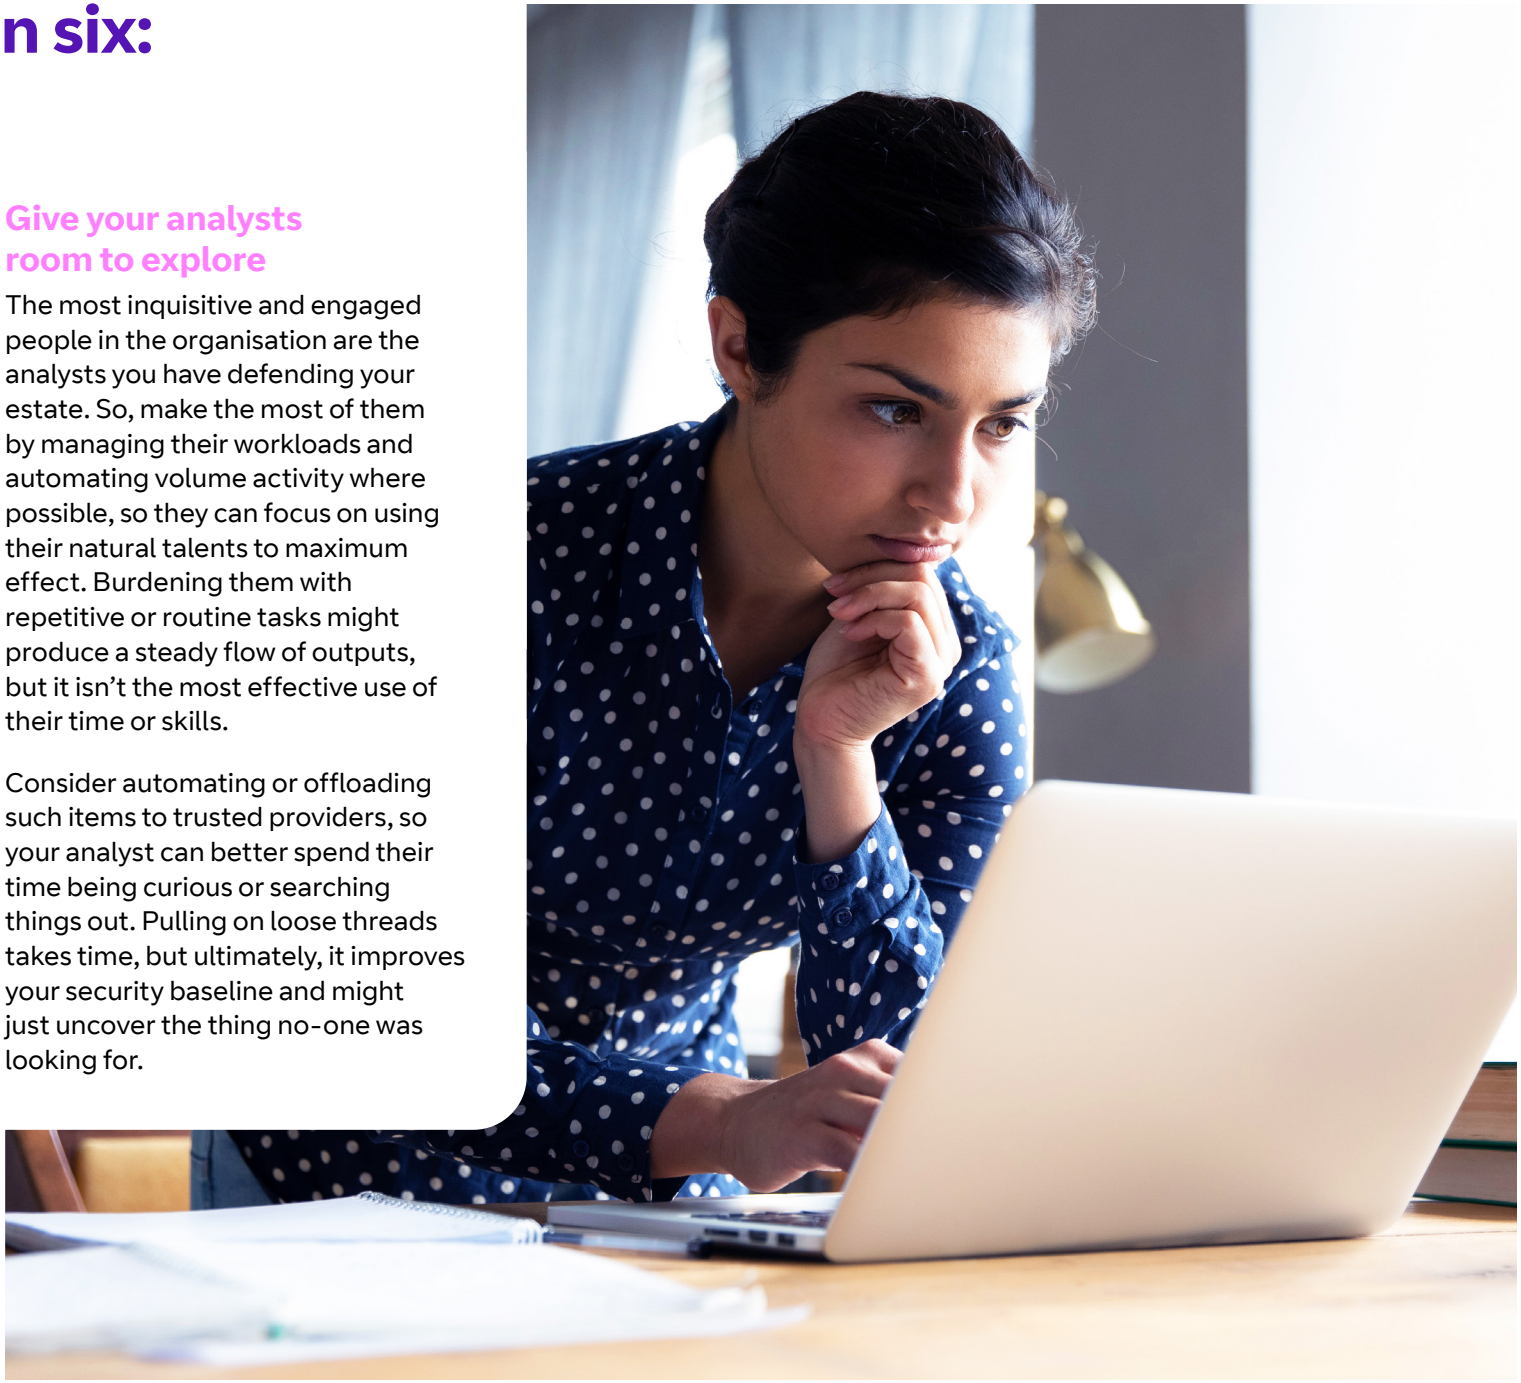

Specific recommendations

Current: protect the infrastructure

Understanding how threat intelligence relates to your organisation is critical. There's a huge amount of widely available information about the attacker behaviours and motivations, but how do you know what's relevant to you?

You'll greatly improve your detection capability if you focus on the tactics, techniques and procedures that are used by attackers, and how their objectives and behaviours could impact your organisation.

You'll also be able to put specific, longer lasting countermeasures in place, as attackers will be forced to change their approach, which costs them time and money. While they grow ever more sophisticated and novel elements of strikes will always occur, it'll always be difficult for them to abandon an entire lifecycle of attack behaviours completely.

Using an industry-approved framework such as Mitre ATT&CK® can help you develop this more strategic approach to reducing cyber risk.

Understanding your attack surface and overlaying the latest threat landscape and current security controls allows you to highlight any gaps in your current strategy.

Configuring your hygiene controls to stop attackers from successfully using typical behaviours increases your opportunities to detect and thwart them. Many cyber-attacks are a question of economics (cost vs. reward), so anything that pushes up cybercriminals' operational costs makes you a less attractive target.

Medium: protect the data

As organisations transform digitally, data (and everything that connects to it) continues to proliferate. And as data grows, so does its attractiveness and value to criminals.

A strong medium-term strategy for dealing with incidents is to build a data protection strategy and to deploy controls to help detect and protect your digital assets.

Our advisory services can help organisations take a data-centric approach to protection, supporting them to define plans for the discovery of information assets, classification, labelling and protection.

We can also partner in defining policy for handling distinct types of data and implement these policies into controls like Cloud Access Security Broker (CASB), web control or dedicated data loss tools.

It's increasingly important to understand what your human, application and machine users are doing within your estate so you can identify any deviations from normal. Our SIEM and SOC services can apply behavioural analysis to telemetry from your IT assets, and our network visibility tools can baseline flows of data and application usage across your estate.

Better, regular visibility of your digital interactions can highlight any anomalous flows of data. And putting a robust data protection strategy in place is critical to making sure your data doesn't flow to those who shouldn't have it.

Future: protect the users

Consider analytics and user behaviour analysis to model normal behaviour so that you're in a better position to detect strange behaviour and deviations. Traditional examples focus on geographical or temporal anomalies but consider first-time activity and unexpected volumes of transactions as well. If you understand how business processes work, you're better able to detect deviations from the norm, and more likely to pick up and prevent fraud and financial crime especially.

What does this mean for the future?

Attackers and criminals will never stop trying to invent new ways of gaining a return on their investments.

If you can make it expensive, difficult and time-consuming for a criminal to achieve their goal, this will limit the range and motivation of attackers targeting your organisation. You need to do the basics to raise the cost of carrying out the crime for attackers and you must accept there'll always be some level of risk to using IT for business advantage.

With this in mind, it's important to keep up to date with the threats and risks facing you, your sector, your employees and your business partners. It's also important to make technology choices that support your future strategy and give you the visibility and control to baseline and measure what's happening.

The identity of people, devices and applications is increasingly critical to how we use IT. Having a thorough understanding of your users' needs and requirements, and adapting to any changes as they happen, is a key part of future Zero Trust strategies. It's also critical to the successful implementation of dynamic automated policy control and decision making.

The 2020 SolarWinds incident shows that, when they're sufficiently motivated, knowledgeable and resourced, cyber attackers can find their way through defences and protections. However, this doesn't mean we should accept defeat.

In fact, it shows how important it is to put a range of complementing security controls and protections in place to provide layers of protection against a dirty network. By doing this, you give yourself the best chance of spotting these attacks early, clearly understanding their impact and scope, and quickly limiting their damage.

Our security services

We're helping customers thrive by delivering world-class security solutions. We have operations in more than 180 countries and support some of the world's largest companies, nation states, and critical national infrastructures. That gives a unique perspective on cybercrime.

Our team of 3,000 security experts in 16 global centres identify and analyse 6,500 potential cyberattacks every day.

The intelligence we gather and evaluate from key sources around the world, helps you spot and deal with the critical threats to protect your business.

We're here to help you keep watch and act decisively at all times.

How to get in touch

Contact a specialist

Offices worldwide

October 2021