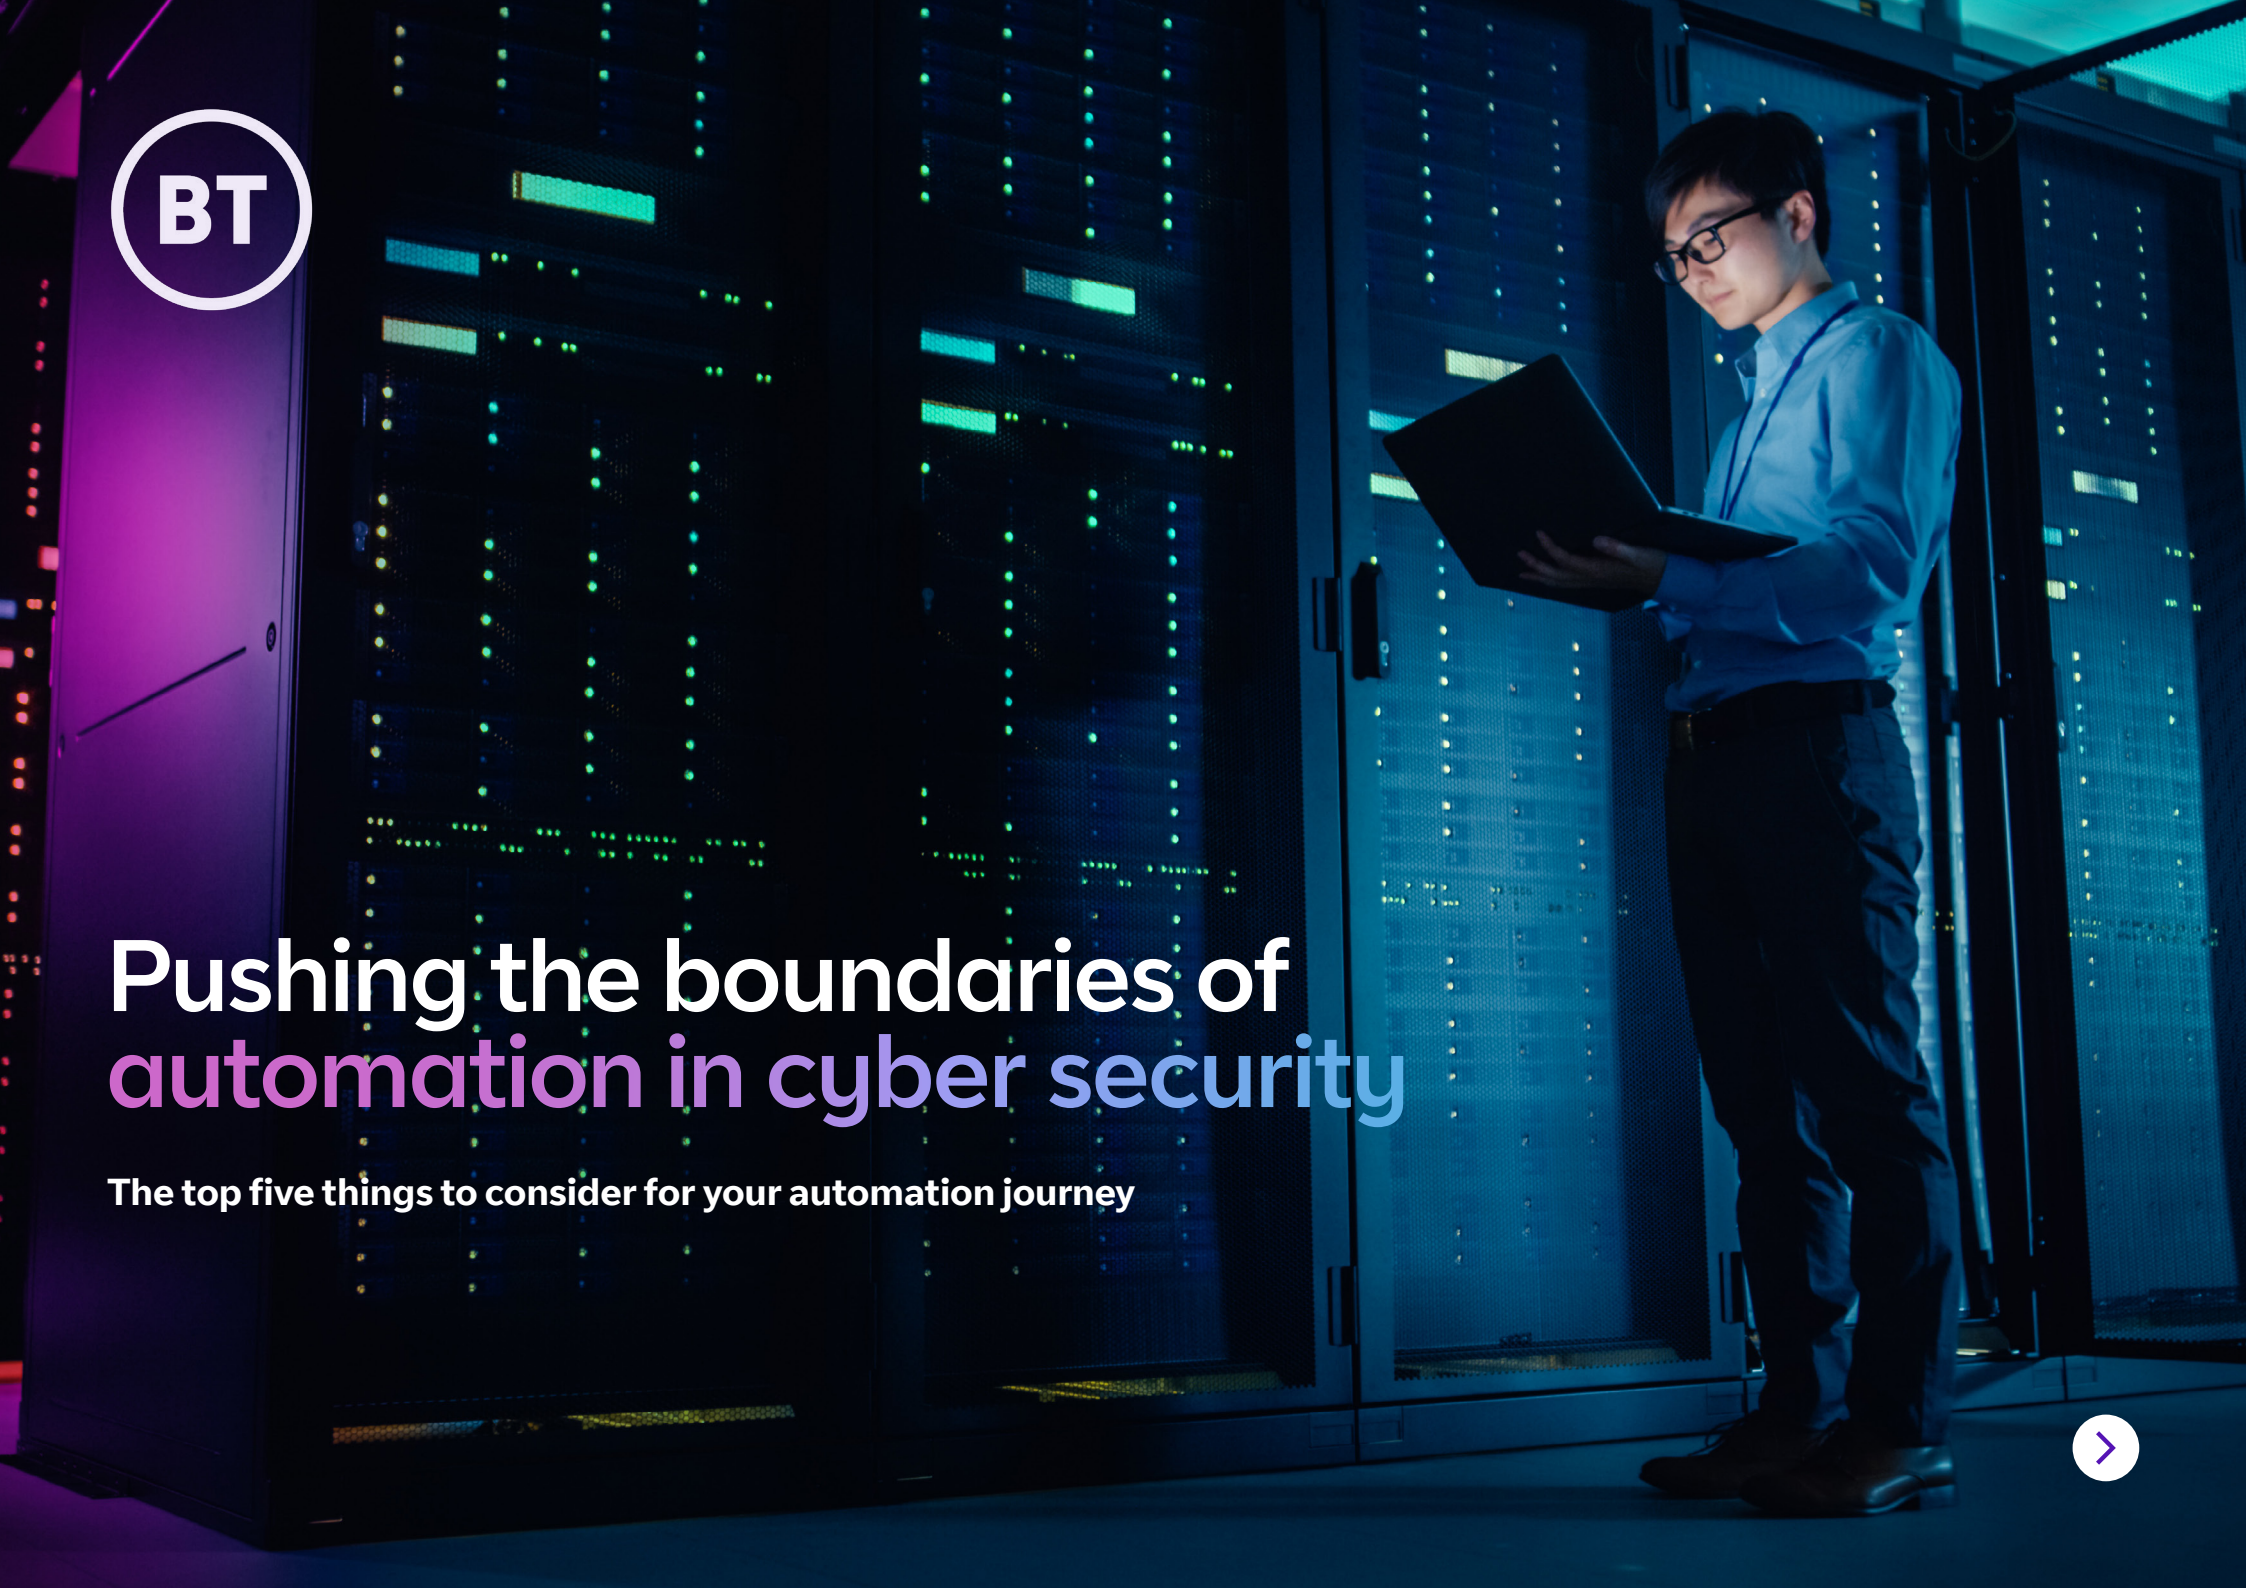

Pushing the boundaries of automation in cyber security

The top five things to consider for your automation journey

Introduction

The sheer scale and pace of today's cyber threats has proved overwhelming for the security department, outpacing its resources. Automation in cyber security is becoming a critical part of an organisation's defences. Now so many modern cyber attacks are heavily automated, enterprises need to adopt more automation to level the playing field.

Automation can help identify relevant intelligence and take on more mundane security tasks. This frees up staff to focus on urgent and high priority threats – something that's especially important considering the shortage of cyber skills affecting so many organisations.

In a recent month, one of our current SIEM clients with 63 active log sources, generated 11.3 billion events. This led to 279 cases for human investigation. Assuming a doubling of log sources, with a corresponding increase in cases, it quickly becomes clear that human analysts are not going to be able to cope.

There's huge excitement and passion around technology's potential to reduce the impact on our people and to optimise our response – but implementation isn't straightforward. We've been using automation to drive business benefit for some years, and in this whitepaper, we'll discuss our own experience and the value it can offer.

We'll also explore the top five things you should consider when embarking on an automation journey. And we'll look to the future, sharing our view on dynamic decision making and the role of humans in cyber response.

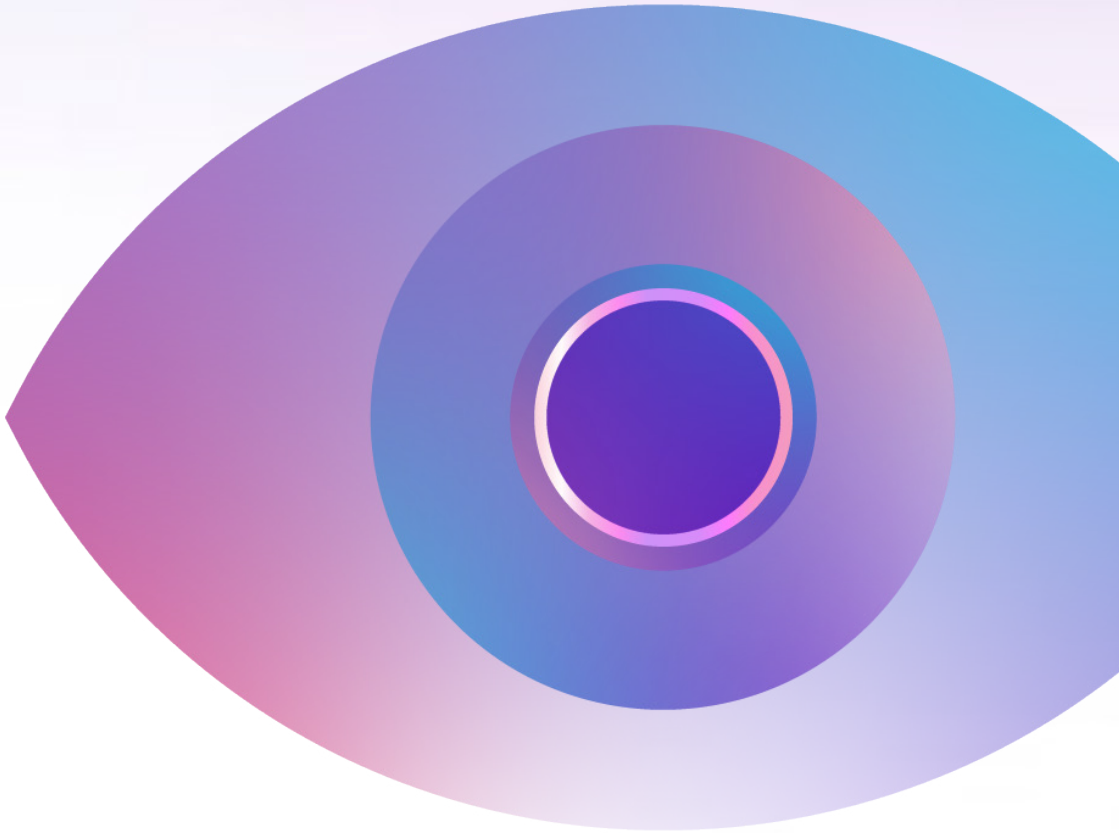

Our automation story

We began the automation journey for our security operations centres (SOCs) in 2018. Our key initial goal was to unify customer experience across all our SOCs. Our mission was to provide great experiences for our customers and, by automating best practice, we intended to improve things still further. Our plan was to provide more consistency by harmonising the customer experience when changes or incidents were handled in different or multiple locations. We also wanted to save time on each change or incident by improving our analysts' efficiency.

However, our initial achievements didn't match our ambitions. Our experience underlined the main learning point we want to share - you should never try to automate a complex process that's not fully documented and well understood. Automating a complex and sometimes uncertain approach or methodology simply leads to a bigger, more complex experience.

The good news is that we found unexpected benefits from a variety of other areas, and we learned from what went wrong, going on to significantly change our approach.

Finding marginal gains in existing processes or automating routine and dull processes delivered real improvements. Yes, this led to resourcing requirement optimisation, but the real benefit was the quality-of-life improvement for our analysts. Satisfaction went up in analyst teams and this helped drive better retention, greater focus and ultimately a better experience for our customers who worked with our analysts.

We've now automated large sections of some of our key playbooks and this has resulted in a more consistent experience. It's saved us significant handling time on many simple service requests and incidents, freeing up our analysts to focus on more critical work. In a few cases, we've been able to significantly reduce the number of different systems analysts we need to use to resolve a situation.

We've learned to pivot to deliver the benefits of automation in two ways:

1. We've made automation a key tenet of our transformational cyber defence platform, Eagle-i. We've built an underlying engine that will enable us to drive improvements for our customers across our portfolio of managed services.
2. We're able to offer more to our large and complex clients who need a dedicated solution. We've applied our initial learning and skills to develop and deliver automation benefits using a few of the leading toolsets, and we have more in the delivery pipeline. Our 'own use' story, coupled with the bespoke solutions and volume delivery via Eagle-i, is allowing us to accelerate the business benefits.

What's next?

The next evolution of our automation journey will see us transform our Security Operations Centres into Cyber Fusion Centres, leveraging advancements in automation, orchestration, Machine Learning (ML) and Artificial Intelligence (AI) to correlate diverse data sets, automate repeatable workflows and prioritise actions for more decisive responses to threats.

We see security automation as critical to ensuring we can deal with the rapidly increasing amount of data, events and incidents generated, as well as delivering a quality threat-focused service for our customers.

Our Cyber Fusion programme is designed to help reduce business risk, cost and activity duplication, and decrease the attack surface, all while improving threat intelligence, productivity and communication. It will also deliver accelerated incident response (reducing Mean Time To Detect and Mean Time To Repair) as well as improved SOC efficiency and output, with capacity freed up to focus on customers' higher value work.

The top five considerations for your automation journey

Organisations that want to learn from our experience and knowledge of the security landscape should consider the following:

1. Skills shortages should drive focus

It's no secret that the cyber security industry has suffered from significant skills shortages for many years. In its 2022 Cyber Security Workforce Study, (ISC)2, a leading international non-profit membership association for information security leaders, estimated a worldwide cyber security skills shortage of 3.4 million people. Despite adding nearly half a million cyber security workers over the past year, the cyber security workforce gap has grown more than twice as much as the workforce, experiencing a 26.2% year-on-year increase. In the face of these supply constraints, the cost of cyber security talent is extremely high, leading to regular churn as skilled employees leave for better pay and conditions. It also makes it very difficult for smaller organisations and less well-funded industries to compete for cyber talent, leading to huge security risks.

Skills shortages dictate workloads

The loss of staff, particularly the more experienced people, has forced those that remain to refocus on the basics and has reduced their capacity to investigate the more complex (and in their eyes, interesting) problems their organisations face.

It is undeniably important to do the basics brilliantly, but there's a cost - staff don't have time to perform root cause analysis,investigate complex new issues and push for continuous improvement and innovation.

Automation can improve working conditions

In an environment where resourcing is limited, it's critical to consider how automation can handle the routine, the repetitive, and those tasks that are important but not urgent. Loading such tasks onto an already stretched team is a great way to demotivate, bore and frustrate them. The result of such a negative workload is often greater attrition - something no organisation wants. Automation offers the ability to process routine and mundane tasks at speed, consistently, and without needing great levels of oversight. It would be naïve to assume that whole tasks can be fully automated, but automating the repetitive parts frees your analysts to apply their skills in more complex ways that widen the breadth of their experience, leading to increased job satisfaction.

2. Look at automation holistically

Automation offers many potential advantages including a reduced reliance on analysts, as well as time and cost savings. However, when looking at automation options for your organisation, you need to assess the overall value of automating each process. The goal is to use automation where it provides the most effective value, without watering down capabilities, introducing additional risk or removing necessary human oversight.

Understand all factors before making decisions

Before implementing automation in your tools be sure you're clear on the benefits and why you're automating. Consider all aspects of your processes, the types of threats you see, talent availability, and the costs / benefits of automating certain decisions. It might make sense to automate highly repetitive daily tasks to free up your analysts' time to focus on large-scale attacks.

You might find that due to the industry you're in you'll need to continue to rely heavily on analysts to make decisions, but that you can use automation to provide your analysts with options that significantly reduce the time to respond. In this scenario, you might choose to use tools that automatically detect a likely threat attack pattern and automatically present your analysts with options to determine the best remediation action for your organisation.

3. Think about the operational implications of automation

Go back just five years and our clients were almost all fearful of systems that could perform automated changes. The idea of giving unfettered access to a 'black box' mechanism that could block applications, systems, or even customers was a little too risky for most. Great for it to suggest or even provide explicit instructions on what to do, but a human needed to be 'in the loop'. For some of our clients, there even needed to be four eyes on each change.

But any human interaction leads to delays, as a human is inevitably slower than an automated process, and this delay increases the risk of something negative happening in the intervening period.

However, almost universally, our clients analysed this risk, weighed it against the risk of loss of availability due to unintended consequences, and decided to keep the human involved. Times have moved on. Other methods of automation have emerged, and our attitudes as a community are changing as computing environments become more software based.

A cloud development example

One of our clients' policies regarding their cloud development environment is relevant here. It's generally well understood that developers and dev environments are one of those places where experimentation is desirable and looser restrictions encourage innovation. This is all good. However, from an attacker's perspective, development environments are the gift that keeps on giving, and in fact have been implicated in many of the larger breaches.

How do we balance these two opposing objectives? In this example, the organisation has a set of security policies around their dev platform that is more open, but also a policy that enforces a mandatory wipe of the entire environment on a monthly basis.

Each and every month it is razed to the ground and redeployed using automated tooling. The security policy is also automated, and it provides a testing ground for some of the riskier or newer controls. Initially, it caused chaos each month, but it's settled into a routine and, amazingly, it all just works.

What the example shows

What's evident from this situation and others, is that clear communication between groups within the organisation is critical to any platform where automated change can happen.

Operational teams, in particular, need to have understanding, sight, and sign-off of such systems to understand the implications. In most organisations, a human 'on the loop' during an initial phase, aware of how such systems operate and able to investigate any change and revert it back in moments, is a wise precaution. After a period of optimisation, tuning and tweaking, confidence will grow to the point that the humans can be 'out of the loop' and the organisation can rely on the automation.

4. Choose your scope and domain wisely

Deciding what areas to automate is a minefield for many organisations. For most clients starting out on the journey, we would recommend focusing initially on basic individual controls and policy enforcement capability. This area tends to be simpler, with less chance of clashes between technology areas and vendors.

Make the most of what you've got

Another area to investigate early on is the automation capabilities built into the tools that you're already using. Many organisations fail to use the tools they've bought to their full extent. Getting the most out of the capability you already have should be an easy first step – both the vendors and your managed services provider will be happy to help.

A much more complex environment is the threat management overlay (SIEM or XDR capability) and any interaction with the control layer to automatically push policy down to multiple controls.

This is complicated due to the number of different tools involved. This complexity means the benefit can be greater, but the risks are commensurately higher. As the number of controls increases, the orchestration demands increase, and so does the possibility that misconfiguration could derail any benefits or the positive perception of any solution.

Use a risk assessment to guide automation plans

Risks also vary by organisation. We've long advocated building a strong understanding of the threats each organisation faces and determining responses based on the tools, techniques and procedures known to be used against them.

This threat intelligence data can be one of the main drivers allowing these responses to be automated – as long as the accuracy of the data is high. In our case, we're putting great emphasis on improving the value we get from our threat intelligence platform and how we share that data around our organisation.

The final related consideration is how you track results. Most measures look at minutes saved per incident, but it's worth considering other data points such as threat intelligence dissemination, incidents detected, threat hunts initiated, and even job satisfaction, to allow you to measure success and focus on continuous improvement.

5. Data clarity drives improvement

Having a clear view of what data is needed, and the insight that it can give you, is key to successfully implementing automation programmes.

Be selective about security data

This is particularly important when linked to security decision making. Security controls are generating more data points than ever before, but before we can use it, we must ingest, process and store all that information. With ever greater volumes of data, the quicker, more accurately and more efficiently a decision can be made, the better the result.

Collection plans and data prioritisation should be central to your data strategy. This ensures that you can consolidate and refine data to reduce volume (and therefore cost) without losing the integrity or value of the data or of the decisions that we can make. Data that's not important to the decision can be an overhead, burning time and / or space, and can cause a delay in reaching a sound outcome.

It's important to identify what critical data you need and where it's located. Inevitably, the data will come in many different forms and from a myriad of locations and sources, and having a well thought through data capture strategy will help deliver the most benefit down the line. With so much data on hand, being clear on validity and prioritisation is fundamental to making sure that automated decision making is accurate and predictable. This predictability is especially important when operating within regulated industries.

Automation into the future

As the world becomes ever more interlinked and connected devices number in the billions, the cyber landscape will continue to increase in complexity. Automation will move from a 'nice to have' into an essential tool that organisations can't cope without.

We will reach a point where the gains of automation are outweighed by the overhead in creating, modifying and updating the processes and tooling – but few organisations are anywhere near this level of saturation. In all such cases, the key to understanding lies in accurate recording of time and impact so as to drive optimum business value.

From a cyber security perspective, we believe that two areas will come to the fore.

The first is dynamic automated decision-making, where systems can react to changing inputs such as context, organisational risk, vulnerabilities etc. and automatically adjust the choices being made. The second is cross-platform remediation, where decisions made by one system can also be automatically implemented in other systems with no human input.

A further consideration is that, while great advances are being made in using automation to help defend against cyber threats, attackers are also deploying automation to try to increase their chances of success. As a result, the importance of contextual, timely and accurate threat intelligence as an input to decision making cannot be overstated.

The role of trust in AI-powered automation

AI-powered automation has been touted as an ideal way to reduce the workload of human analysts, improve the speed and accuracy of detection and response, and enhance the overall resilience of security operations. However, AI also poses some challenges, such as how to ensure the quality and reliability of automated processes, how to mitigate any potential biases, hallucinations and errors of AI models, how to have trust and certainty with AI-powered automation processes, and how to achieve operational certainty and compliance in a dynamic and unproven environment.

Trust is a vital factor for cyber security automation, as it affects the acceptance and adoption of AI solutions by human users. Trust can be influenced by various factors, such as the transparency, explainability, accountability, reliability, and usability of AI-powered automation. To build trust, cyber security practitioners need to adopt an analyst-centric approach that addresses all these factors.

- **Transparency** - providing clear and accessible information about the goals, methods, assumptions, limitations, outcomes, and impacts of AI-powered automation.
- **Explainability** - providing understandable and meaningful explanations about how and why AI-powered automation works and behaves in different situations.
- **Accountability** - providing mechanisms for identifying and attributing the responsibilities and liabilities of AI-powered automation and its developers.
- **Reliability** - providing guarantees for the consistency, accuracy, robustness, and repeatability of AI-powered automation.
- **Usability** - finding ways for AI-powered automation to be easily incorporated into our existing systems and processes and to support analysts in their time-consuming tasks.

Pushing the boundaries of automation in cyber security

Our research labs at Adastral Park in Suffolk are using AI and ML extensively to create autonomy in prevention, detection, and response capabilities. For advanced detection, our research is focused on autonomously learning from data. Here, we use advanced deep learning and other AI technologies to learn from data where possible and, when data isn't available, learn from human decision-making. Human-in-the-loop (humans are part of the training and decision-making process) and human-on-the-loop (humans act in a supervisory role) are important research differentiators in our programme.

In the response layer, our focus is on automated and semi-automated response capabilities. Our vision is to have combined proactive and predictive capabilities in our network to determine and orchestrate response to an attack throughout the network to create a 'self-healing' capability. Our research is focused on epidemiological modelling of threat propagation and re-enforcement learning to determine the most appropriate response to an attack at any point in a kill chain.

As we develop our security operations centres into Cyber Fusion Centres, we're looking to incorporate the best of existing ML and AI discoveries and usage from diverse parts of our business, industry and research into our cyber security automation. We recognise that how exactly this capability will apply in an operational environment will need further discussion and development. However, we've already identified some key issues to be address.

- **Certainty of the course of action** - how confident can we be to allow an AI-powered automation to act?
- **Motivation and intent** - can human analysts discuss and rationalise their approach and options?
- **Compliance** - do we have the ability to prove to an independent body, auditor or regulator the predictability and repeatability of an AI-powered automation's actions?

> AI-powered automation is an exciting new area. The bold and brave will press on, and doubtless they'll encounter many issues which we'll need to discuss, deliberate and agree on before we see more widespread adoption.

Our Security Advisory Services

Our security advisory team is here to support you on your automation journey. We offer strategic security guidance and solutions to organisations across the globe, reflecting the market demand for expert guidance to navigate today's complex cyber security landscape. Our advisory services help organisations at all stages of their security journey to assess and test their defences and select the solutions that match their security needs - whether that requires building an entirely new security strategy or upgrading their protections to combat the latest threats and trends.

We offer a range of advisory services to help our customers with their security strategies.

Security assessments

Governance, risk, compliance

- Threat prioritisation
- Cyber maturity assessment
- Cyber risk quantification

Strategy and optimisation advisory services

Evaluate, define, recommend

- Transformation strategies
- Security roadmaps
- Framework architectures

Offensive security

Evaluate security defences

- Penetration testing
- Assume breach
- Red teaming

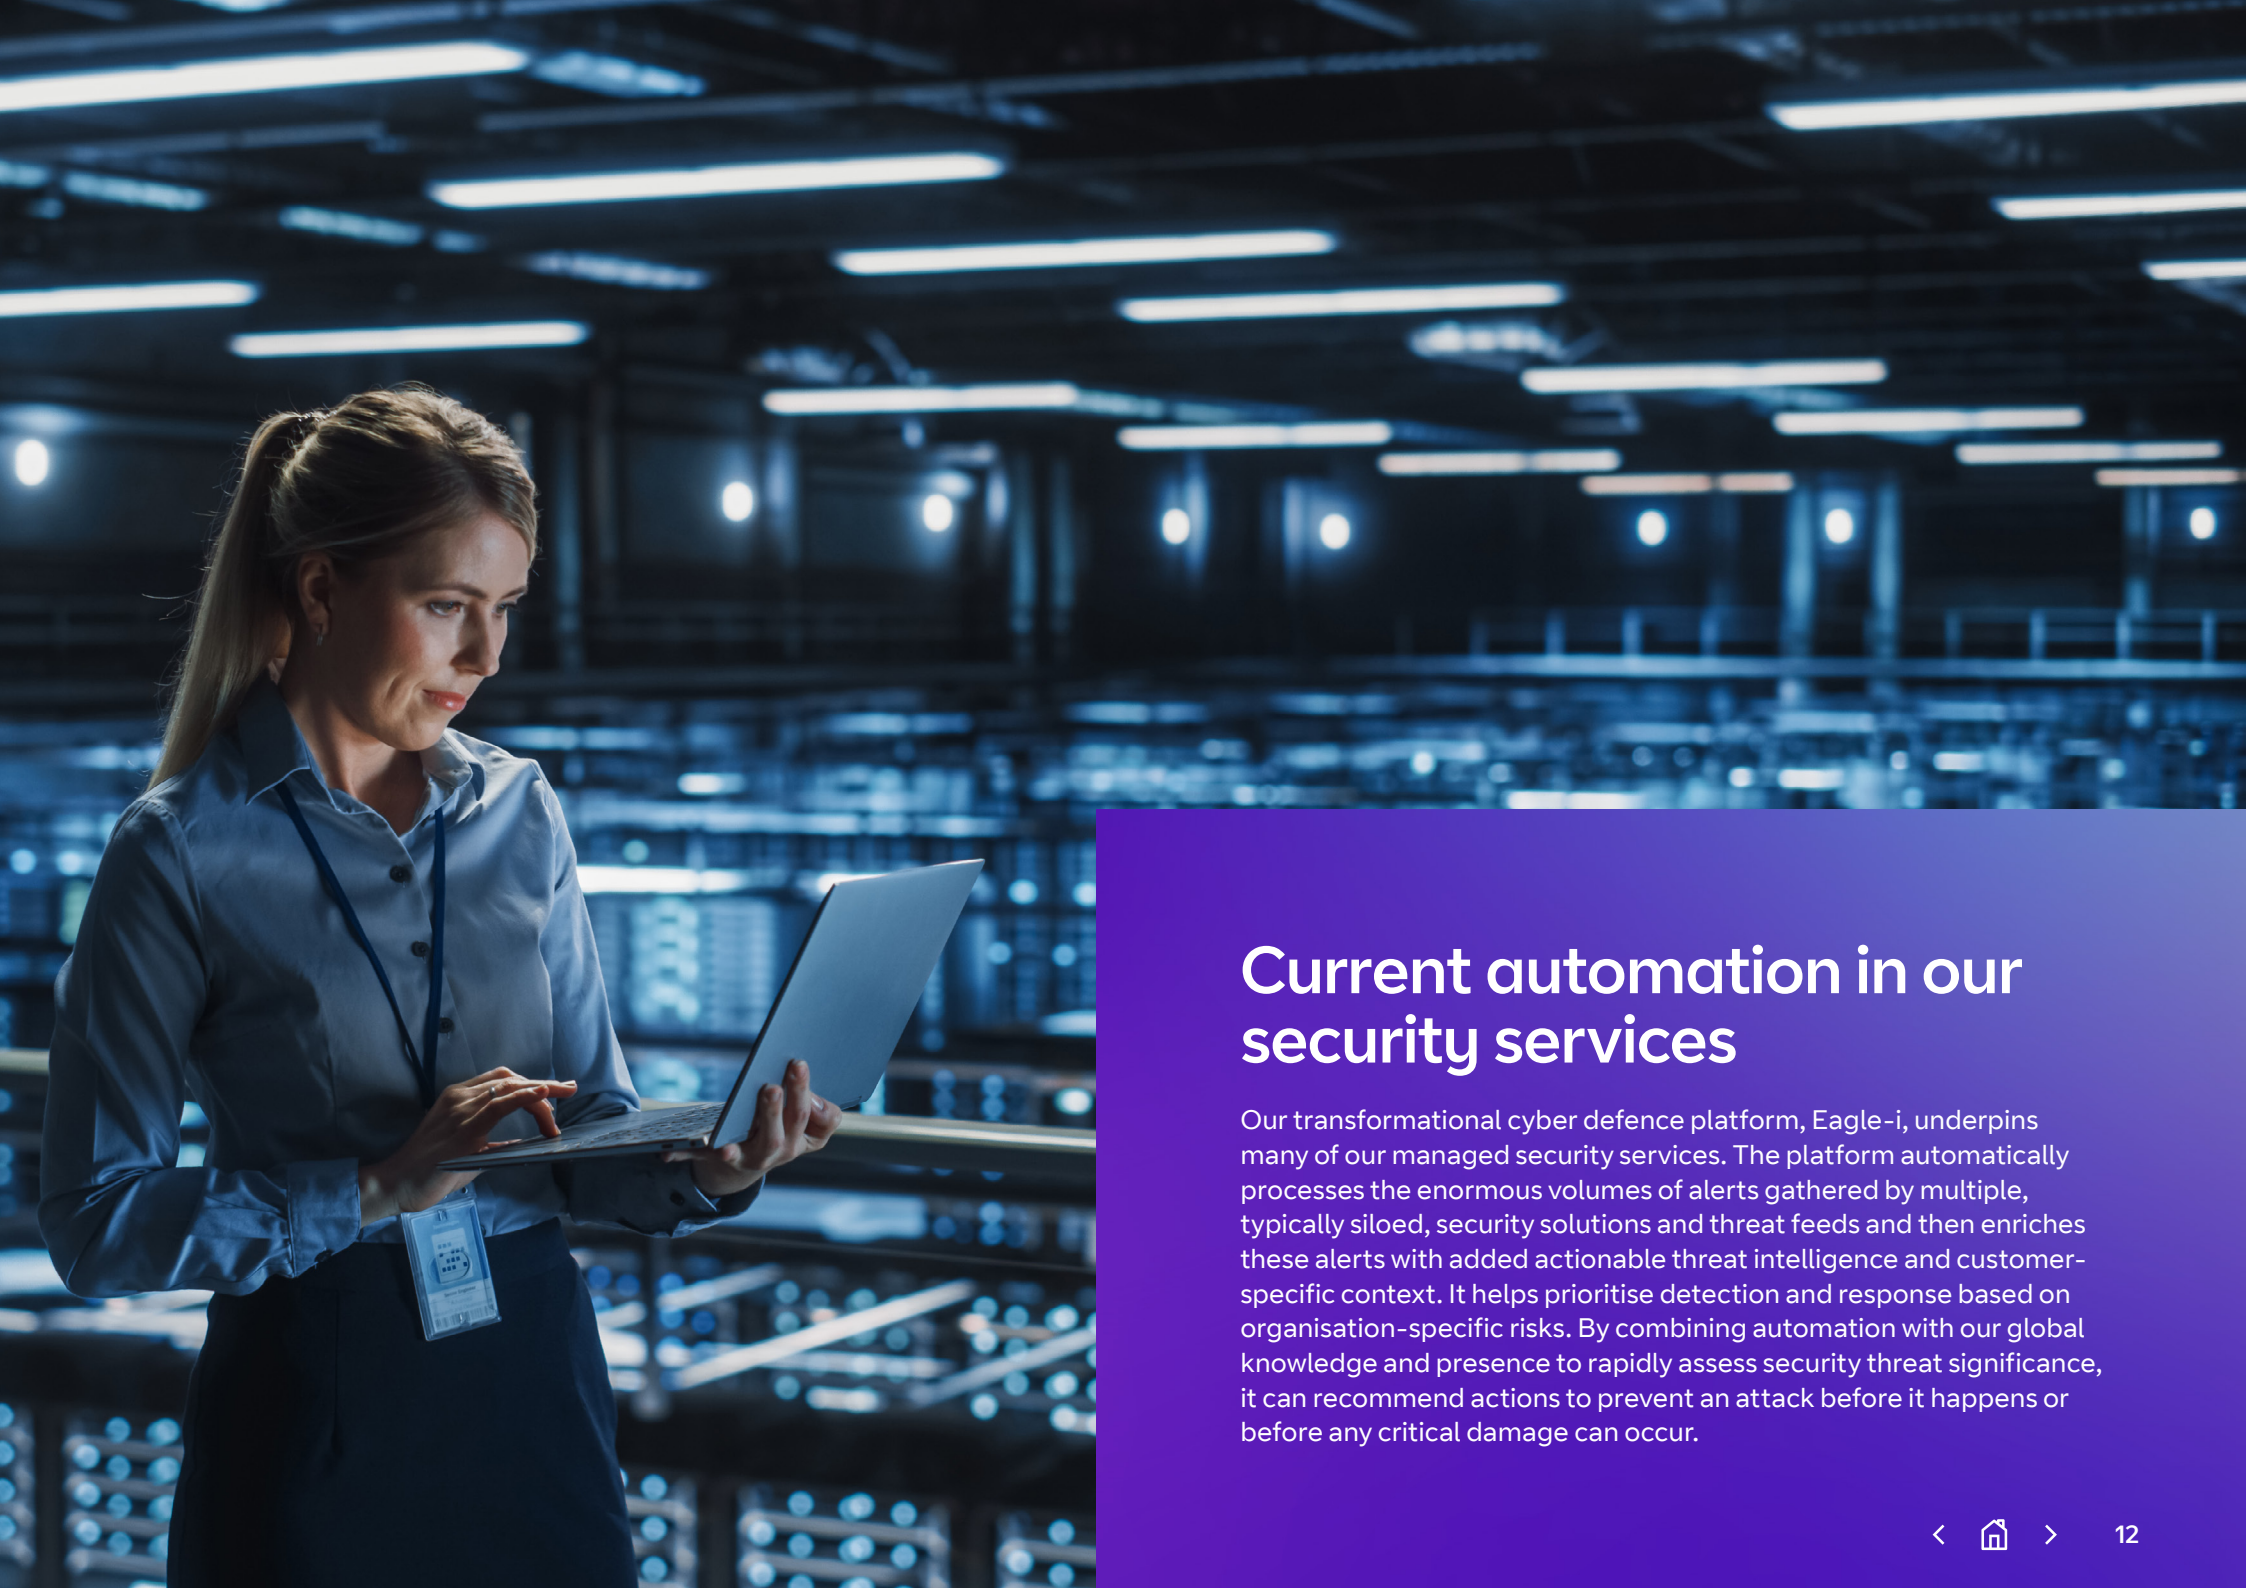

Current automation in our security services

Our transformational cyber defence platform, Eagle-i, underpins many of our managed security services. The platform automatically processes the enormous volumes of alerts gathered by multiple, typically siloed, security solutions and threat feeds and then enriches these alerts with added actionable threat intelligence and customer-specific context. It helps prioritise detection and response based on organisation-specific risks. By combining automation with our global knowledge and presence to rapidly assess security threat significance, it can recommend actions to prevent an attack before it happens or before any critical damage can occur.

Develop your security automation strategy

We're ready to help you explore and evolve. Start your journey by finding out more about our security capabilities: www.bt.com/security

Offices worldwide

The services described in this publication are subject to availability and may be modified from time to time. Services and equipment are provided subject to British Telecommunications plc's respective standard conditions of contract. Nothing in this publication forms any part of any contract.

October 2023