SYGNIA

OCTOBER 2024

CROWDSTRIKE FALLOUT: NAVIGATING THE RISKS OF INTRUSIVE SECURITY TOOLS

Learn how to balance robust cybersecurity with operational stability in the wake of the CrowdStrike outage.

KOBY ZVIRSH, DONOVAN PACE, HAIM NACHMIAS

Contents

EXECUTIVE SUMMARY

In July 2024, a CrowdStrike update led to widespread outages, raising concerns about the risks associated with intrusive security tools in critical systems. This incident serves as a reminder of the delicate trade-offs between security and operational stability.

Intrusive security tools, while essential for defending against advanced cyber threats, can compromise system performance, particularly in operational technology (OT) environments where uptime is crucial.

To address this challenge, Sygnia recommends several strategies, including robust testing, deployment and disaster recovery processes, tailoring security controls to different assets, applying defense-in-depth, fostering close collaboration.

This article highlights the need to continuously adapt to strike the right balance between security intrusiveness and operational stability, ensuring resilience in an ever-evolving threat landscape.

INTRODUCTION

Intrusive security measures, which require deep integration into IT systems to monitor, control and defend against threats, have become a standard approach to safeguarding business-critical IT infrastructure from both routine and unforeseen risks. However, the recent widespread outage caused by CrowdStrike's popular endpoint protection solution has reignited concerns about the hidden costs and potential risks associated with these measures. This incident, which disrupted thousands of organizations worldwide, has prompted a reassessment of the delicate balance between security and operational stability, resurfacing long-standing dilemmas regarding the trade-offs involved in using intrusive security tools.

As technology advances and the incentive for threat actors grows, the complexity and scope of security threats have also expanded, necessitating more comprehensive security measures. But while these tools are crucial for achieving resilience against cyberattacks, the intrusive control they can exert over IT systems can lead to performance issues or even system failures, causing frustration for IT teams tasked with ensuring operational availability. Although designed to protect critical infrastructure, intrusive security tools can sometimes undermine the very stability they are meant to ensure.

The adoption of any type of intrusive security measure in a business-critical environment necessitates a thorough examination of our technological framework, understanding the trade-offs, and tailoring the approach to our specific operational and cyber risk tolerance. This strategic approach ensures that we are well-prepared for unforeseen incidents and helps prevent security solutions from turning into liabilities.

UNDERSTANDING THE CROWDSTRIKE INCIDENT

Before delving into the broader risks of intrusive security measures, let's first examine the most recent incident to trigger this conversation. On July 19, 2024, a content update involving Channel File 291 for CrowdStrike's Falcon sensor software led to significant system disruptions globally, by causing the infamous "Blue Screen of Death" (BSOD). The update was intended to improve threat detection through Named Pipes, a Windows feature for inter-process communication. However, an error in the update file's formatting caused the CrowdStrike driver – operating in kernel mode, and thus having unrestricted access to system memory – to attempt to access a non-existent memory address, triggering the BSOD.

The BSOD issue affected a wide range of industries. For example, airlines canceled or delayed flights, healthcare providers experienced issues with medical records, and up to 100 emergency 911 call centers in the US reported downtime. The outage disrupted businesses that depend on operational technology (OT) as well as those using conventional IT systems, further highlighting the widespread dependency on reliable technological services across various sectors.

CrowdStrike swiftly identified and addressed the issue, but the solution required customers to boot systems into Safe Mode or the Windows Recovery Environment, resulting in extended downtime. The challenge was even greater for remote endpoints, as they required physical intervention by an operator, further delaying recovery efforts.

This scenario is not unique in the realm of intrusive IT software. Historical examples of malfunctioning cybersecurity tools include a 2010 McAfee update that misidentified a critical Windows system file as malware, a 2016 Symantec update that caused BSODs on Windows XP machines, and a 2019 Trend Micro update that led to crashes and network disruptions due to compatibility issues with certain Windows versions.

As with these past cases, the CrowdStrike mishap not only revealed vulnerabilities in the security software and its supply chain processes, but also demonstrated the fragility of modern IT systems and how interconnected they are. This incident serves as a stark reminder of the potential disruptions that highly complex and invasive security tools can cause when they malfunction, despite their crucial role in defense strategies.

CHALLENGES AND CONSIDERATIONS OF INVASIVE SECURITY TOOLS IN CRITICAL ENVIRONMENTS

The deployment of intrusive security tools such as Endpoint Detection and Response (EDR) systems, Network Firewalls, Network Access Control (NAC) solutions, Intrusion Prevention Systems (IPS), Web Application Firewalls (WAFs), and Security Orchestration, Automation and Response (SOAR) platforms, plays a pivotal role in maintaining business continuity and safeguarding IT systems. These tools, which combine monitoring, prevention and automated response, are essential in mitigating threats in real-time and defending against evolving cyberattacks. Similarly, patch management, while not a tool, is an intrusive process essential for mitigating security vulnerabilities.

These solutions often go beyond basic and static prevention, offering advanced analytics and automation capabilities to detect attacks based on behavioral anomalies on top of known threat signatures. In dynamic and high-stakes environments, where threat patterns are constantly shifting, these capabilities are critical for identifying previously unknown threats. However, they can also be intrusive, as they require deep integration with IT systems and the authority to intervene in real-time.

These control capabilities can lead to inadvertent operational disruptions when these tools trigger false positives, suffer from software issues, or have unintended configurations. For example, an EDR system might mistakenly quarantine a critical process, halting production in an industrial environment or delaying real-time financial transactions. If it has kernel-level privileges, as with most EDR systems, it can even cause system-wide crashes that cannot be readily reverted, as occurred in the recent CrowdStrike incident.

As the threat landscape continues to evolve, the need for these intrusive security tools becomes more apparent. However, greater reliance on such tools increases the risk of operational interference. The consequences of these disruptions can range from minor inconveniences to catastrophic failures, depending on the criticality of the systems involved. In sectors like healthcare, finance, or industrial operations, where uptime is crucial, even a minor disruption can lead to cascading failures, further highlighting the delicate balance between ensuring robust security and maintaining uninterrupted operations.

Specific Challenges in Operational Technology (OT) Environments

In OT environments, where system availability, integrity and safety are paramount, deploying security solutions can present significant operational risks. The Purdue model, which has become the de facto standard for segregating OT networks, defines several layers of control systems and computing infrastructure, each with distinct characteristics and roles in the control process. These characteristics influence the feasibility of implementing security controls at these layers. At the lower network levels, such as Levels 0 and 1, which encompass sensors, actuators, and controllers that directly interact with physical processes in Industrial Control Systems (ICS), deploying security tools is typically not feasible due to lack of hardware and software compatibility. On Levels 2 and 3, which are responsible for controlling the automated control process, security controls are feasible but pose significant operational risks, with the primary concern of disrupting operational stability and affecting system safety.

The potential for severe consequences at these Purdue model levels makes organizations reluctant to deploy invasive tools like EDRs, which can lead to system conflicts or performance degradation. Additionally, OT software vendors often refuse to support customers who use unapproved invasive tools like modern EDRs, allowing only specific anti-malware solutions that have been thoroughly tested and vetted. These simpler security measures are often preferred by organizations, as they have minimal impact on system performance and are considered better suited to work with legacy systems. They can provide a basic level of protection from threats such as malware introduced via USB sticks or unauthorized software, while maintaining operational stability and safety.

As we move up the Purdue model to Levels 3.5 (Industrial DMZ) and Level 4, where systems manage plant operations and enterprise IT infrastructure, the environment begins to resemble traditional IT settings. Here, the emphasis typically shifts toward cybersecurity, and the deployment of more sophisticated and invasive security solutions becomes more accepted and necessary.

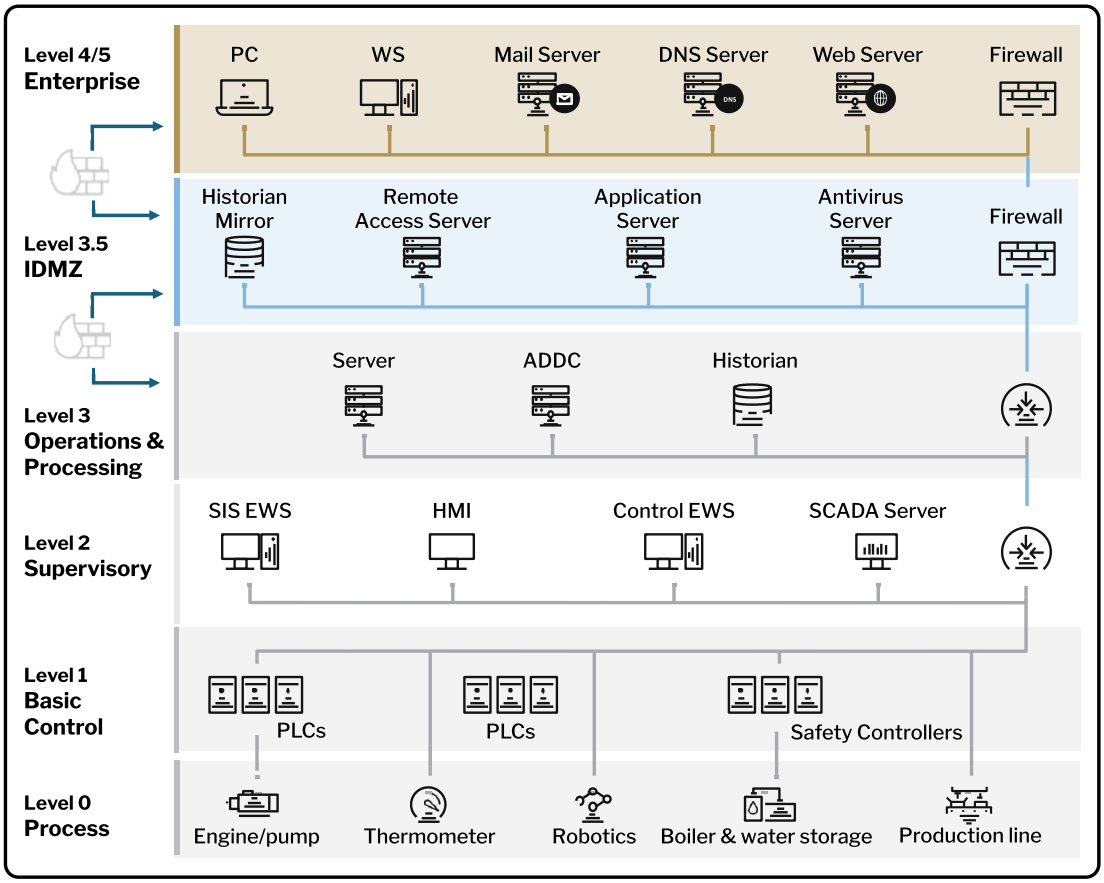

Purdue Model - Reference architecture for industrial enterprise

This shift is driven by the increased complexity of the threat landscape and the growing exposure to outside systems through the corporate network and the public internet. The potential risk of physical damage at these higher levels is lower, while cyber risks begin to outweigh concerns regarding system interference. This makes it feasible to introduce more advanced tools that enhance security without compromising the overall integrity of the OT environment.

MINIMIZING THE RISK: BALANCING CYBERSECURITY WITH BUSINESS OPERATIONS

In today's digital landscape, not all organizations look the same. Some move fast and adopt modern technologies, while others depend on more traditional systems. This diversity means that security strategies must account for the specific needs of each business, balancing the priorities of efficiency with the need for robust cybersecurity. This chapter explores how organizations can minimize risk by aligning their cybersecurity efforts with business goals, ensuring both protection and continuity.

Proactive Risk Management

Identify key assets, assess risks – including availability risks from operational failures – and implement appropriate protective measures.

The first step in minimizing cyber risks is identifying the organization's key assets, including where sensitive data is stored, which systems process it, and which processes support essential business functions. Once identified, it is essential to evaluate their exposure and associated vulnerabilities, to ultimately assess the potential impact and likelihood of exploitation. Security teams typically assess these assets through the lens of the CIA triad – confidentiality, integrity, and availability. While confidentiality risks primarily stem from cyberattacks, integrity and availability may also be jeopardized due to operational failures.

Proactive risk management involves implementing protective measures before threats materialize, reducing the risk of both cyberattacks and operational disruptions. By their nature, security teams lean towards the most restrictive controls. However, to be effective business partners, they must understand the potential effect on business operations. While certain controls may enhance protection, they might also introduce complexities or slow down processes. By weighing the benefits of controls against their potential impact on business activities, security teams can ensure that key assets are well-protected without compromising the organization's ability to function effectively.

An additional consideration is third-party vendor risk. Relying heavily on a single vendor for critical business functions can introduce concentration risk, where a failure in one external service or tool could cascade through the organization, leading to widespread operational disruption. Mapping dependencies across vendors and assessing risks beyond cybersecurity, such as operational resilience, helps mitigate these risks.

Managing Software Updates

Use pre-deployment testing, staged rollouts, and rollback plans to minimize disruptions.

Would you tolerate a 5-minute downtime when your smartphone updates itself? While your answer may instinctively be “yes”, it will probably not be so if it occurs during an important business call. For individuals and organizations, timing and context set the rules.

Every software update or process adjustment carries the potential to disrupt system availability. The recent CrowdStrike incident demonstrates the consequences of failing to detect a faulty update, leading to widespread disruptions. To mitigate these risks, organizations must adopt effective change control strategies. These strategies help ensure that issues, when they arise, are caught early, preventing them from escalating into significant operational disruptions.

Pre-deployment testing combined with a staged rollout is one of the most effective ways to minimize the risks associated with software updates. This method starts by testing updates in a controlled environment, then deploying them to pilot groups (also known as deployment rings) with lower operational risks. These groups help assess the update against a predefined success criteria before it is rolled out to a broader or more sensitive environment. Although keeping systems updated with the latest versions of Indicator of Compromise (IOCs, often labelled as ‘content updates’ by vendors) and of software is typically preferred for security reasons, organizations may choose to somewhat delay updates when system stability is a priority. This can be achieved manually or through appropriated product features, offered by some of the vendors.

Even with the best testing and rollout strategies, unforeseen issues can still arise. That’s why having an effective rollback plan is essential, which would allow to quickly revert to a previous stable version if an update introduces unexpected problems. This plan should be well-documented and tested, ensuring it can be executed efficiently under pressure and minimize downtime.

Throughout the update process, maintaining clear communication with all stakeholders is critical to success. This includes not only IT and security teams but also business units and end-users who may be affected by the update. Close collaboration ensures that all parties are aware of the update schedule, potential risks, and any steps they may need to take. Stakeholder alignment fosters collaboration, minimizes confusion, and builds trust across the organization, contributing to a smoother update process.

Defense in Depth

Implement layered security controls to strengthen protection and reduce reliance on any single tool.

As security professionals, we know that relying on a single line of defense is a common pitfall. Leveraging a Defense in Depth approach not only allows for stronger overall defense against threats through layered security measures but also helps to better balance risk with operations. By implementing multiple controls to address vulnerabilities at different levels, security teams can reduce reliance on any single, potentially intrusive tool. For example, combining a less intrusive monitoring solution with robust perimeter defenses and stringent access controls can result in effective overall security posture without significantly impacting system performance.

Consider a scenario where a company employs Defense in Depth by using a combination of network firewalls, operating system hardening, strong authentication and least privileged access control. Even if one layer, such as a network firewall, is compromised or becomes less effective, the remaining layers continue to provide protection. This layered approach can help minimize the need for more intrusive tools that could disrupt operations, offering a balanced and resilient security strategy that aligns with business objectives.

Expect the Unexpected

Develop robust disaster recovery and business continuity plans, test them regularly, and foster antifragility by using disruptions as opportunities to strengthen and adapt.

The days of backup robots and tape libraries are (largely) behind us, replaced by cloud providers offering five-nines SLAs. Yet, despite these advancements, organizations still struggle to recover from unexpected disruptions. A disaster recovery plan (DRP) that is not integrated with the broader operational aspects of a business continuity plan (BCP) may prove ineffective in practice. The same goes for simplified recovery drills, which may miss critical cross-functional steps likely to be needed in real-world scenarios.

The organizational BCP and DRP strategies must be tailored to the business. Ensuring confidentiality and ensuring availability require different approaches. Identifying critical business processes, dependencies and potential single points of failure, while learning from past incidents and from other organizations in the business sector, can provide valuable insights into what could go wrong. For example, the CrowdStrike incident highlighted the risks of relying solely on remote recovery tools, as physical intervention was required for some systems. This underlines the importance of accounting for both physical and remote recovery options, where appropriate.

It is incredibly challenging to predict all of these factors accurately the first time. For this reason, organizations should practice comprehensive simulations, identify missed components, and learn from them. Over time, continuous testing and real-world disruptions will not only improve recovery readiness but also help the organization become antifragile – emerging stronger and more adaptable after each challenge.

Promoting Positive Collaboration Between IT and Cybersecurity Teams

Building strong collaboration between IT and security teams is key to balancing security needs with operational performance.

There is inherent tension between IT and security teams. Each team plays pivotal role in maintaining the ongoing operations, reliability and security of the company. IT teams are often focused on optimizing technology for productivity. The security team, on the other hand, is dedicated to safeguarding the organization's digital assets and protecting them from unauthorized access, data breaches, and other cyber threats. The collaboration between these two teams is essential to balance usability and security.

By working together, IT and security teams can mitigate the challenges of deploying potentially invasive security tools. For example, rather than imposing a one-size-fits-all security tool that might slow down system performance, the teams can collaborate to develop a more refined approach, one that enhances security while preserving system availability and performance.

Establishing a healthy, collaborative working environment takes time but should be viewed as a long-term investment. This can be achieved through frequent and open communication, which builds trust and transparency. Acknowledge each team's objectives, be open to hearing other opinions, learn from failures, give credit, and celebrate collective successes – all of which will strengthen collaboration and support meeting mutual goals.

STRATEGIES FOR OPTIMIZING SECURITY WITH MINIMAL DISRUPTIONS

As cyberattacks become more sophisticated, security teams must adopt strategies that strengthen protection while minimizing disruptions. This section focuses on practical methods for deploying security controls that maintain system performance and align with operational needs, helping organizations stay resilient and efficient as security challenges continue to evolve.

Tailoring Security Policies to Assets

Customize security controls based on the asset's role, sensitivity, and exposure to ensure effective protection without unnecessary interference.

Not all assets require the same security measures. By understanding the specific role and risks associated with each asset, organizations can apply controls that are both effective and proportionate. For instance, internet-facing assets demand stronger security measures, while internal assets with less exposure may rely on less disruptive methods, reduces excessive interference.

As another example, servers and endpoints require different considerations. Virtual servers in high-availability clusters can tolerate individual failures and be quickly restored, making them better suited for more intrusive security measures. In contrast, physical endpoints are typically less easily restored. For corporate user workstations, intrusive security measures remain crucial due to direct threats such as phishing and web browsing, and because their downtime is typically manageable. However, remote OT endpoints may be both less exposed and less tolerant of downtime, requiring less intrusive security solutions.

Security policies should also reflect the criticality and sensitivity of each asset. Systems handling sensitive data might prioritize confidentiality, accepting potential downtime for stronger protections. Meanwhile, assets critical to operational continuity require measures that maintain availability without disrupting performance. Notably, even systems seemingly less critical can be exploited by threat actors for lateral movement toward more sensitive assets, making it essential to apply robust security also to assets that may serve as 'the weakest link' as part of attack scenarios in organizational networks.

Detective vs. Preventive Controls

Balance preventive and detective controls based on the risk tolerance and operational priorities of your environment.

Typically, the ability to prevent an active attack requires the capability to detect it. However, predicting the exact consequence of a security control can often be challenging. For instance, this unpredictability might stem from a highly dynamic environment where processes go up and down constantly, or from a newly deployed tool that the team is not yet fully familiar with. Sometimes, diligent fine-tuning of the control is necessary to ensure it works effectively without causing negative interference.

In network environments or situations where the risk of system disruption cannot be tolerated, intrusive preventive tools and false positives pose too great a risk. In such cases, detective controls offer a viable alternative. Examples of such controls can be IPS or WAF systems configured in monitoring mode, allowing them to detect and alert without actively blocking traffic. Likewise, an EDR can be configured to prevent only certain types or levels of threats while only issuing warnings for others. This balance helps minimize the impact on system performance while still providing valuable threat visibility.

Keep in mind that available resources are crucial when implementing detective controls. A high volume of alerts can overwhelm security teams, leading to missed incidents if there are not enough personnel to manage them effectively.

Decisions about when to favor preventive or detective controls should be made on a case-by-case risk-based approach. For example, if the risk of an unblocked event would be too high – such as on a senior executive's workstation – accepting some false positives may be necessary to ensure greater protection.

Operation Models of Security Tools

Consider the operation models – such as inline, parallel, agent-based, agentless, kernel-space, or user-space – when selecting security tools.

When deploying security solutions, it's important to understand their various operation models and their implications on system performance and stability. The key question is whether the level of intrusion is justified by the security benefits, within the business context.

In the context of network security tools, such as Intrusion Detection Systems (IDS) and Network Detection and Response (NDR), inline systems analyze and control traffic directly on the data path. This can slow down overall system performance due to the volume of data being processed. Parallel configurations, on the other hand, operate alongside the data path, scanning a mirrored copy of the traffic. This approach minimizes direct impact on system performance, although it may not be as effective in real-time threat mitigation.

Agent-based solutions delve deeply into system operations for comprehensive monitoring, often at the cost of higher resource consumption which can lead to performance bottlenecks. Agentless solutions, though less resource-intensive, offer limited visibility and may miss crucial interactions that only a deeply embedded agent could detect.

When assessing endpoint protection agents, the distinction between tools operating in the kernel space of the operating system, such as CrowdStrike Falcon, and those functioning only in user space, is significant. Kernel-mode operations allow for deep system integration, offering robust monitoring and intervention capabilities but carry a higher risk of system crashes due to the level of access required. Agents operating in user mode, while safer and less disruptive, might not catch deeply embedded threats as effectively and could be more easily bypassed by threat actors.

Considering the risks, Microsoft has previously attempted to limit third-party kernel-level access, notably in 2006 with Windows Vista, but faced resistance from cybersecurity vendors and regulators. The CrowdStrike outage has brought these concerns back to the forefront, leading Microsoft to hold a security summit to discuss reducing reliance on kernel-level access and improving the resilience of the Windows ecosystem, while ensuring security tools remain effective. Organizations should monitor these developments, as future changes could impact decision making.

Network Security

Segment networks and use advanced firewall features to limit exposure while minimizing potential performance impact.

A well-designed network architecture can help protect the organization's systems while minimizing intrusiveness. This becomes even more important in environments with limited capabilities to run agents on servers, workstations, or other equipment, such as in OT networks, particularly in their layers 0-2 of the previously described Purdue model.

A fundamental control is network segmentation which limits access to endpoints and servers by dividing the network into isolated segments. The more isolated the environment, the smaller the attack surface and blast radius in case of a breach. This approach reduces the risk of unauthorized access and the spread of threats within the network.

Modern firewalls include advanced capabilities such as IPS and network-based antivirus. These can provide additional layer of protection for the organization's assets, especially where endpoint-based controls may be limited.

There are additional techniques that can be employed, depending on their availability in the product. One such technique is controlling a "fail" condition. For environments where the risk of downtime is too great, certain tools can be set to a "fail open" state, ensuring that they do not block traffic if they encounter issues. This configuration may be preferable in certain high-risk scenarios compared to a "fail close" configuration.

Minimizing Attack Surface and Exposure

Reduce the attack surface by hardening systems and limiting entry points to lower the likelihood of exploitation.

Minimizing the attack surface reduces the number of entry points available for exploitation, substantially lowering the risk of a breach. This approach can also serve as a compensating control when intrusive security tools are not feasible. One effective strategy is the use of hardened or minimized operating systems out of the box, such as Red Hat Enterprise Linux (RHEL) with Security-Enhanced Linux (SELinux) or Ubuntu Core, a lightweight, security-optimized version of Ubuntu. These systems are pre-configured to limit vulnerabilities, reducing the need for additional, potentially disruptive, security measures.

Hardening can also include restricting unnecessary services, enforcing least privilege, and enabling security features like full-disk encryption, secure boot, or Microsoft Defender Credential Guard. These methods bolster security without adding significant operational overhead.

Application control is another relatively non-intrusive way to reduce attack surface, by ensuring that only trusted software can run on a system. If implemented using an allowlisting approach, it can reduce the risk of malware and unauthorized applications without the need for constant, intrusive scanning. This method is particularly effective for static workloads, which are not expected to change frequently. For dynamic workloads, a blocklisting approach can be employed, although it may not provide the same level of security as with allowlisting.

Out-of-band File Scanning

Out-of-band scanning can be used in sensitive environments to detect malware without impacting system performance.

Anti-virus technologies are an old and proven method for detecting and protecting against malicious code. However, in certain environments, such as OT or production servers, deploying such solutions may be challenging either due to lack of official support from the manufacturer, or concerns about performance degradation. When this occurs, compensative scanning tools can be used to mitigate risk.

Security teams need to analyze file transfer flows and design the controls accordingly. For example, in an isolated OT environment, file transfer between networks may occur using USB devices. Implementing a process that ensures each USB device is scanned before use can significantly reduce the risk of introducing malware into the system. In cloud environments, scanning server storage outside of the running OS can be used to detect malicious files without impacting performance.

While these solutions may have security drawbacks compared to the use of local antivirus and advanced EDRs, such as the lack of real-time prevention or of memory visibility, they can minimize the risk of infection and limit the attack surface without causing performance issues in sensitive environments.

CONCLUSION

The recent CrowdStrike incident serves as a stark reminder of the potential risks associated with deploying highly intrusive security tools. While these tools are vital for detecting and mitigating sophisticated threats, they also pose risks to operational stability. As business processes continue to undergo digital transformation and reliance on third parties grows, the next global outage is not a question of if, but when. Whether it is caused by a security tool or another factor, organizations must be prepared.

In today's constantly evolving threat landscape, vigilance and continuous adaptation are essential. Key strategies are robust testing, deployment and disaster recovery plans, tailoring security controls to match assets' risk profiles, reducing the attack surface, and wisely applying defense in depth. Together with others, these approaches help strike the right balance between security and operational efficiency, ensuring resilience in the face of evolving threats.