WithSecure™ Countercept presents

The Moment of Truth

Stories from the eye of the cyber storm

Mehmet Surmeli

Senior Incident Response Consultant

John Rogers

Senior Incident Response Consultant

Joani Green

Managing Consultant, Incident Response

James Dorgan

Senior Threat Hunter, Continuous Improvement Lead

WITH secure

The moment of truth

Despite heavy investment in cyber security, organizations continue to lose the battle against attackers – who continue to adapt. The business impact of a successful attack is often significant, and highly public, often with further consequences for wider society, representing a chilling moment of truth for an organization. This shouldn't be the case.

Security controls are failing because organizations don't respond at the right moment. Response should shut down an attack as it occurs, preventing business impact; all too often, the response comes too late, in the form of a slow, costly recovery from a successful attack.

By lavishing attention on detection and recovery rather than detection and actual response, organizations miss a vital opportunity to respond before material business impact occurs.

The good news is that much of this impact can be averted by moving the point at which response begins much closer to the moment of detection. A decisive First Response slashes exposure to the very real business risk of a high-impact cyber security incident.

In this paper, we'll explain why response delays exist, why mitigation of business impact, not speed of response, should be the core measure of success, and why dynamic approaches that combine timing, human expertise and strong technology constitute what we call a strong First Response.

This report draws on the real-world experiences of experts from WithSecure's Incident Response and Detection and Response Teams.

Changing motivations have resulted in more attacks – and caused more damage

Recent years have seen a fundamental shift in attackers' approaches, leading to an increase in the volume and impact of attacks. In turn, this has exposed security controls used by many organizations – from small businesses to major enterprises – as wholly inadequate.

One of the biggest drivers behind this shift is money.

Success encourages imitation, and high-profile, extortion-based cyber security attacks can generate significant returns for attackers. The number of threat actors involved in these types of attacks has consequently skyrocketed and attracted attackers previously focused on espionage. Verizon noted[9] in its 2020 Data Breach Investigations Report (F-Secure is a contributor) that 86% of the breaches on which it collected data were financially motivated and a tenth were espionage-related. The 2021 edition of the same report shows an upward trend from there.

Not just an enterprise problem

A consequence of threat actors shifting towards financially motivated extortion attacks is that it broadens the potential victim list; not every organization holds intellectual property rights or highly classified information, but very few organizations can operate without risk of extortion.

Put another way, this is not a problem that affects only cash-rich, high-profile organizations.

A report from Coveware identified that the median employee number for businesses hit by ransomware attacks in the last quarter of 2020 was just 234, with 35.7% of ransomware attacks affecting organizations with an employee count of between 101 to 1,000.

Security through obscurity is not an effective strategy for mitigating these types of attacks.

Criminals don't just profit from successful extortions

Extortion is not the only revenue stream in the cyber crime ecosystem. Threat actors now frequently claim to have exfiltrated data while deploying ransomware in order to raise the chances of a payout. Public data leaks risk embarrassment, regulatory fines or reputational damage, creating an incentive for victims to shut up and pay up.

This double extortion model is relatively new and relatively well-known, but attackers don't have to perpetrate the extortion to make a healthy profit.

How? Although the extortion of victims is the endgame, attackers are happy to buy and sell individual elements of an attack. Much as the licit business world industrializes, over time threat actors have specialized, creating what is sometimes referred to as the 'cyber crime ecosystem'.

Initial Access Brokers – Threat actors who gain initial access to organizations via phishing, compromising RDP accounts or exploiting software vulnerabilities. They then sell this access to other actors who are responsible for the extortion itself.

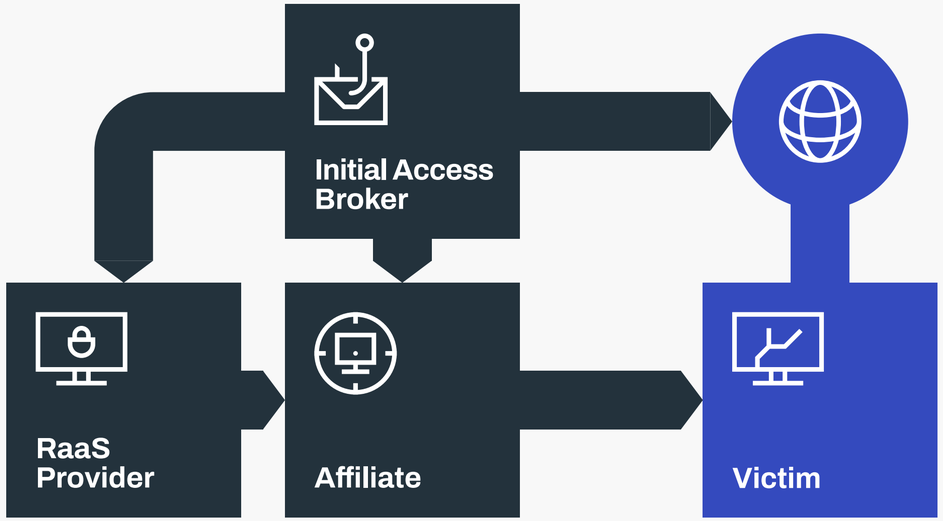

Ransomware-as-a-Service (RaaS)Providers – Platforms that provide threat actors a suite of features required to successfully extort an organization; e.g., packaged malware and attacker tooling, training services, victim communication channels and cryptocurrency laundering. The providers generate revenue by charging a license fee to use the platform, or by taking a cut of the ransom generated from a successful attack.

Affiliates – Threat actors responsible for further compromising the victim organization and extorting them through the deployment of ransomware and exfiltration of data. They scale their efforts, and therefore their revenue stream, by using the services provided by Initial Access Brokers and RaaS Providers.

Profit among thieves

For example, Initial Access Brokers are paid by other attackers for bundles of accesses, Ransomware-as-a-Service (RaaS) Providers offer their platform and services at another stage (for a consideration), and so on. The best suppliers, as in the licit world, attract more work at a higher rate.

All of this means that, even if you have no intention of paying a ransom and have every expectation that you can respond to and recover from an attack, you're still a target for someone who will be paid regardless.

One more thing

These two factors may have increased the potency and volume of attacks, but another factor has also contributed to the rise of ransomware. Businesses have opened systems, applications and data to the cloud, giving online access to staff members, partners and customers from a multitude of devices and locations. This has vastly increased the number of potential entry points to be protected.

Put simply: the attack surface many organizations present to the outside world has increased in size exponentially, and attackers do not have to complete their attack to make money.

All hope is not lost

There's always a moment of truth, during which it's possible to fight back and win.

Based on our own experiences and the structure of the cyber crime ecosystem itself, we know that extortion-based attacks follow a lifecycle where ransomware deployment does not immediately follow initial access. Threat actors want to maximize their chances of a payday by identifying and deploying ransomware on systems where downtime will really cause some pain – encrypting the files of the first machine they gain access to isn't going to be enough leverage to convince the organization to pay the ransom.

The moment in time between identifying an initial access and deployment of ransomware on business-critical systems varies from hours to days (sometimes referred to as the 'time to objective'). This creates a window of opportunity to detect and remove an attacker **before** they have achieved their objective and caused material impact on the organization.

Although this moment is almost always present, the reality is that organizations often can't, or don't, make use of this golden opportunity: their systems and incident response retainers are set up to recover, not necessarily to respond.

Recovery isn’t response

What many organizations would define as ‘response’, we would define as ‘recovery’.

Often our Incident Response team will be called in after an attacker has deployed ransomware or severely compromised an organization’s systems. At this point, our Incident Response team’s efforts are focused on bringing the organization back online, acting in an advisory capacity to senior management and conducting post-breach analysis to determine the root cause. These are all necessary and valued components of recovery – not response.

We know our Incident Response professionals are worth their weight in gold, yet they’re also often the first to say they’d rather not have to be in a position to answer calls from organizations in trouble. One common refrain from the Responders who contributed to this report was that they felt they were often talking to people on the very worst days of their careers. Helping those individuals avoid Groundhog Day by building the right response approach is something our Responders placed great value on.

Response costs less than recovery

The anatomy of current extortion-based attacks provides a window of time where an attacker can be detected and removed before they have caused material business impact.

Removing an attacker after they’ve gained initial access but before they’ve done serious damage comes at a low business cost. An hour to remove the malware or half a day to reimage the machine is a trivial cost compared with recovering from a full-blown incident. Even if the attacker can compromise multiple machines before they are removed, the cost to business is still limited to a days’ worth of remediation effort.

In comparison, the cost to business if the attacker achieves their objective is significantly higher. While the remediation effort alone is much higher (weeks and months in our experience), the real cost to business is the loss of revenue and potential regulatory fines. Cognizant estimated their loss in revenue for a quarter to be between $50-$70 million[16].

If **responding before** an attacker has achieved their objective is much cheaper than **recovery after** an attacker has achieved their objective, why isn't everyone doing it?

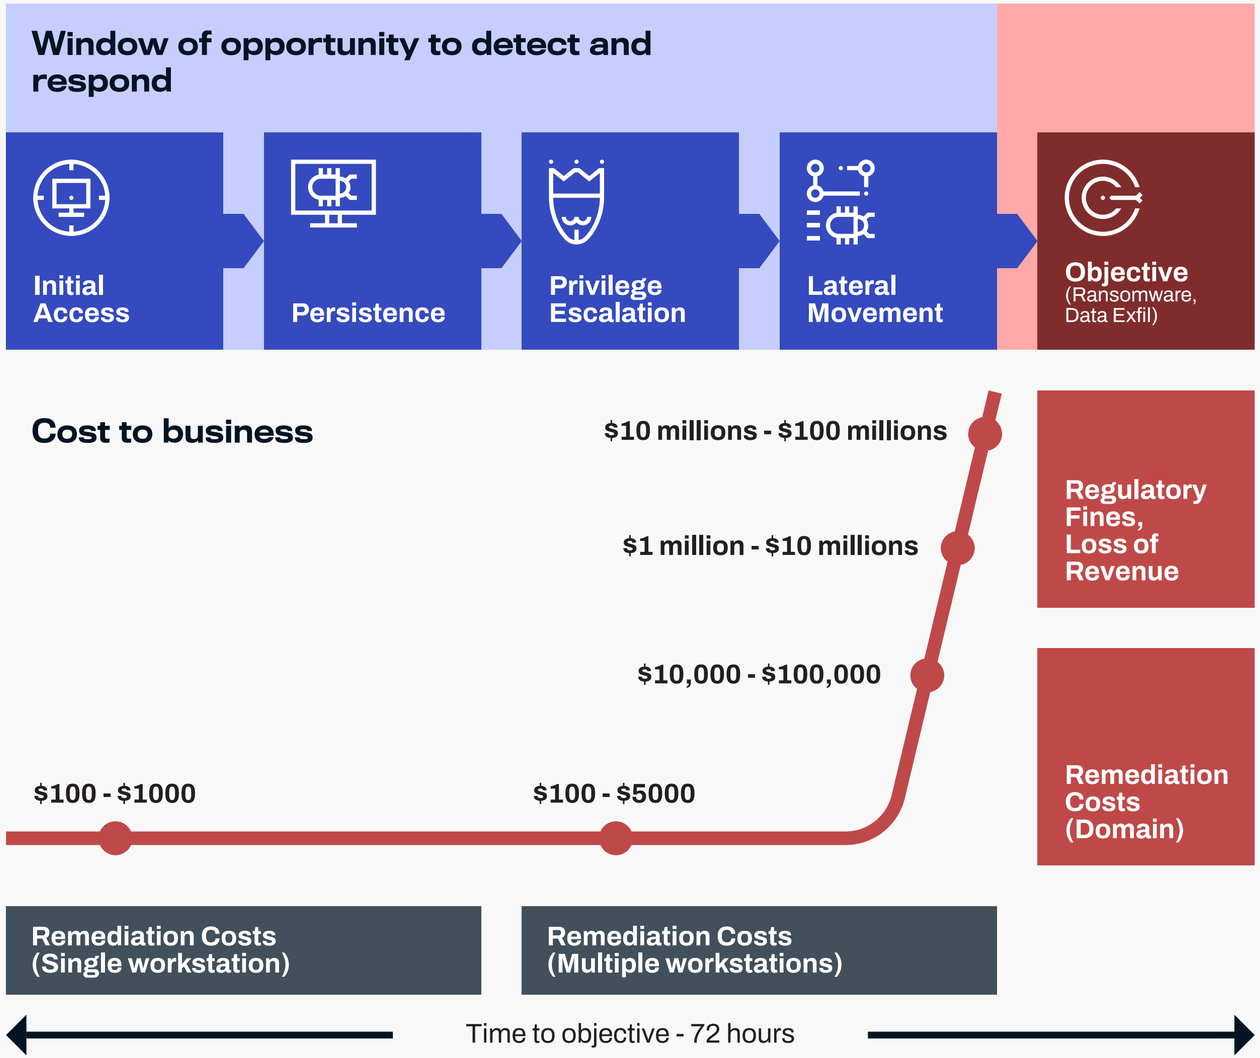

Why does the Response Gap exist?

Organizations that lack the capability to respond to an attack before it causes business impact have what we define as a Response Gap, a concept we have discussed before. A Response Gap can exist for many reasons. Some of the main ones observed by our Incident Response and Detection and Response Teams are as follows:

Reason	Example
Organization has detection technology but no one to monitor its output.	Victim's antivirus detected commodity malware, but lack of personnel to quarantine the file led to a full compromise of their server estate.
Personnel do not have adequate response training or experience, leading to ineffective responses.	Victim with Endpoint Detection and Response received multiple critical alerts and immediately isolated those machines from the network, leading the attacker to reactively deploy ransomware on hundreds of other machines that the victim had not identified as being compromised.
Existing security controls are not configured to be 'response-ready'.	Victim's gateway devices had minimal logging enabled, hampering Incident Response efforts to investigate and increasing the amount of time the attacker was on the network.
Lack of budgetary preparation leading to a delay in Incident Response starting their investigation.	Walk-in victim took a week to get a Scope of Work signed, turning a commodity malware infection into a partial domain compromise.

That the gap exists is not necessarily the fault of the organization, it's more a reflection of a change in the environment. Adapting to new circumstances is, however, a necessary step. The approach we'd suggest is something called First Response.

First Response is effective response

We have talked a lot about response being the removal of an attacker before they achieve their objective. How does this work in reality?

The intent of First Response is to identify the root cause and extent of an attack before initiating a containment plan that will remove the attacker from the network before they cause business impact.

The key principle of First Response is making sure you determine the extent of the attack before starting containment. An effective containment is one where the attacker is eradicated in a single act. If you know where the attacker is, you can do this.

A rushed containment, one where the extent of the attack is not understood before it starts, can often lead to **unintended consequences.** These include the attacker reactively deploying ransomware on any machines they still have access to, or the attacker going back into stealth mode, resulting in a drawn out (and costly) game of cat and mouse. Unintended consequences do not give the response outcome you are looking for.

There's always a trade-off to be made at this point. It's sometimes easy to rush to conclusions on an attack and initiate a plan that turns out to be half-baked. Automated responses, or those initiated from playbooks, often rely on speed rather than effectiveness – something which can trigger consequences from an attacker that's only been partially evicted. It might seem counterintuitive for a response methodology to advocate slow rather fast response, but an effective response requires a blend of speed and precision based on solid reconnaissance.

First Response makes use of the window of time between initial compromise and ransomware deployment. First Responders map the approach and extent of the potential breach, picking up details, patterns and indications that would go unnoticed by an automated response to the first chirp of an alarm. Speed is one thing, but measurement of velocity also has to include a measurement of effectiveness: sitting on the fastest train in the world is a thrill until you notice there are no brakes installed.

The key measurement after First Response should be whether the incident resulted in business impact after the event.

A practitioners' view of what makes a good First Response

Our Incident Response and Detection and Response Teams identified five key characteristics of an effective First Response:

1 Visibility

Being able to draw strong conclusions on how an attacker accessed and compromised a network is vital for devising and executing a containment plan with a high degree of confidence. Visibility starts with an Endpoint Detection and Response agent that captures a broad set of data and forensic artifacts.

2 Attacker Knowledge

Having experience and knowledge of an attacker's modus operandi contributes towards an effective containment plan. For example[19], knowing the typical working hours of an attacker helps identify the best window for execution of the containment.

3 Stakeholder Management

Attacks are high-pressure and uncomfortable situations, particularly for individuals experiencing it for the first time. Keeping stakeholders calm by explaining the process upfront, keeping them informed and being available to address concerns is vital.

4 Integrated Tooling

Having the detection, investigation and response features available in a single tool minimizes the delay and human error that could be the difference between a successful and failed containment plan.

5 Human Involvement

Ensuring human beings are on hand to bring context, experience and empathy to bear is important. Tooling and automation can be counterproductive if they're not backed up by human experience and intuition.

First Response in action

To illustrate what we mean by First Response, we'll tell the stories of two incidents that demonstrate the effectiveness of this methodology.

Case study

1: A carefully formulated response during initial deployment

Our team was starting deployment of WithSecure's Countercept Managed Detection and Response (MDR) service at a large power equipment manufacturer. Our experts immediately detected the presence of Cobalt Strike, a commonly used attacker framework. Further investigation found 11 infected machines, including six domain controllers. Techniques used by the attacker were consistent with a ransomware gang we'd previously dealt with, giving the team a good indication of what the attacker's next step would be, and when it would happen. In this instance, having a high degree of confidence that the attacker was likely going to deploy ransomware at scale over the weekend gave our team a good idea of how much time they and the client had to develop a response.

Had the threat gone undiscovered, or had the organization tried to eject the hackers before the extent of the breach had been confirmed, the attack could have halted production at the company for an extended period and prevented the company's ability to trade.

Working closely with the customer's security team, we were able to develop a remediation plan that could be implemented quickly and effectively without the risk of alerting the hackers. Together, we killed the malicious processes while the customer blocked access at the firewall and reset domain credentials, removing the attacker in a single act and preventing it from returning.

Key points:

1. We were able to develop and execute a containment plan quickly, despite being at a potential disadvantage: we had discovered the attack in its late stages as we rolled out our agent to a new customer. We could do this because we had dealt with the attacker before and understood how it operated.
2. Cost to business was significantly reduced because MDR, with First Response, was deployed in time to catch the attack before material business impact took place.

2: Traditional detection and recovery vs. First Response

Perhaps the best illustration of why correct response is critical is our next story involving a large organization without effective response in place and one of its subsidiaries – an WithSecure™ Countercept customer.

The parent organization was running Microsoft's E5 Suite, employed a regional Managed Security Service Provider and retained an Incident Response service from a large consultancy. Attackers managed to bypass the organization's defenses and, once inside, obtain elevated privileges. This access enabled them to breach the systems of the subsidiary company – our customer.

The parent organization engaged their Incident Response provider, who deployed tens of consultants for an extended period, costing the company a six-figure sum. The organization was told it had been the victim of an advanced attack, possibly by a nation-state actor.

Over-reliance on tech-driven responses may look good on paper but can be dangerously weak in practice if not combined with dynamic input from experienced response teams.

Case study

Parent company

Security posture: MSSP, SIEM, Email Security, EDR, AV

Subsidiary company

Security posture: Purpose built MDR

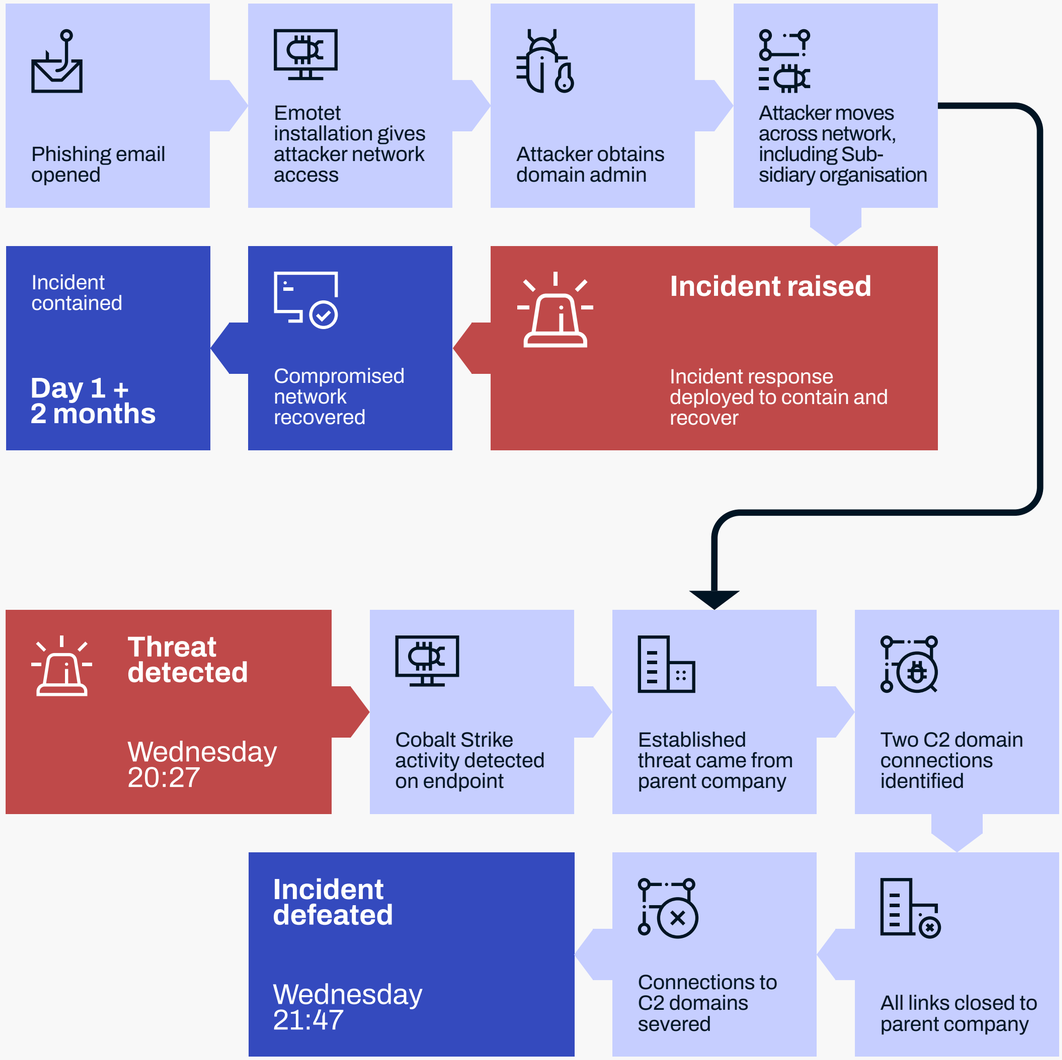

A phishing email is never ‘game over’

When we conducted our own post-incident analysis, we discovered that the attack had not been an advanced nation-state actor at all but an opportunistic ransomware gang. One of the parent company’s users had opened a phishing email that downloaded Emotet, a commodity banking trojan-turned-botnet used to gain and sell initial access. Given its prevalence, this should have been detected and removed from the organization well before the attackers gained the elevated position they then used to infiltrate the subsidiary.

In terms of cost to business and disruption, the contrast between the two response strategies could not be more stark: one hit six figures in recovery costs alone, while the other was resolved as part of retained cost, with little-to-no business disruption.

Conclusion

Moving response teams and technology as close to the point of detection as possible has gone from being desirable to mandatory.

Detection and response must be prompt, so it's understandable that many security vendors talk about the rapidity of automated or push-button detection and response. We don't think this is the best approach: humans need to be involved. It's critical that the team making the First Response is closely aligned to any Incident Response team, and responses need to be comprehensive rather than automated and incomplete.

Difficult as it may be in the midst of responding to a potential incident, the measure of success needs to be the business impact after a response, not the speed at which the first shot is fired. A kneejerk reaction often results in more damage than a careful, planned, comprehensive and overwhelming eviction of the invader.

Good First Response demands a combination of human expertise and purpose-built, integrated tooling placed close to initial points of compromise, allowing First Responders to quickly understand what's going on during an incident and shut down potential compromises.

If a vendor can take ownership of an organization's detection and response in one go, it can be truly effective for its customers. But for many organizations, this can be a pretty big leap of faith. The provider has to be aligned to existing processes. Access and responsibility can't be handed over lock, stock and barrel – that level of trust and confidence must be earned.

Strong MDR services that help identify, detect and respond to attacks before they become incidents with a decisive First Response makes a massive difference for many organizations. The skills of Detection and Response Team members, especially when coupled with dedicated tools and support, allow organizations to avoid turning a drama into a fully fledged crisis.

To better understand your organization's ability to respond decisively to incidents, rather than falling into the trap of recovering from successful attacks, take our three-minute test at https://www.withsecure.com/detect-to-respond, and speak to a member of the WithSecure™ MDR team to find out more about how First Response can relieve business impact and risk for you.

Who We Are

WithSecure™ is cyber security's reliable partner. IT service providers, MSSPs and businesses along with the largest financial institutions, manufacturers, and thousands of the world's most advanced communications and technology providers trust us for outcome-based cyber security that protects and enables their operations. Our AI-driven protection secures endpoints and cloud collaboration, and our intelligent detection & response is powered by experts who identify business risks by proactively hunting for threats and confronting live attacks. Our consultants partner with enterprises and tech challengers to build resilience through evidence-based security advice. With more than 30 years of experience in building technology that meets business objectives, we've built our portfolio to grow with our partners through flexible commercial models.

WithSecure™ is part of F-Secure Corporation, founded in 1988, and listed on the NASDAQ OMX Helsinki Ltd.

WITH secure