Gallagher Re

Looking from the Outside-In: Can taking the threat actors' viewpoint help insurers?

Looking from the Outside-In: Can taking the threat actors' viewpoint help insurers?

The last two years have seen events occur that were almost unimaginable in the modern age; a global widespread pandemic and the invasion of a sovereign country in Europe. Aside from the terrible human cost and impact on day-to-day lives, the events have also sent shockwaves through the Insurance industry. This was observed firstly, with COVID where global lockdowns of entire countries were a previously unforeseen (and unpriced) possibility, and more recently with Russia's invasion of Ukraine. COVID has caused 'PML bust' events for the most heavily impacted books like Event Contingency, whilst the Aviation market is facing potentially similar outcomes from the 10bn+ worth of insured trapped assets on Russian soil (at the time of writing). The balance between premiums charged for these risks vs the potential aggregate losses results in a payback period well in excess of 100 years. This is clearly unsustainable for the market and requires a fundamental change in the way in which risk is assessed, measured and priced for, particularly for poorly modelled perils.

Another line of business that intuitively has the potential to cause widespread, systemic losses is Cyber. Whilst still a relatively small class compared to the total insurance industry, there are a multitude of reports that describe scenarios in which a Cyber event could cause insured losses in excess of $200 billion or more. It is clear that regulatory scrutiny on the class is increasing, and shareholders/boards will be putting pressure on firms to robustly evaluate their systemic exposures (including Cyber) after recent events. This may lead to further updates in pricing tools, reviews of internal capital models and deployment of capital.

The past two years have seen the Cyber product navigate a difficult time in its growth as attritional ransomware and increased loss ratios forced heavy rate rises which, in turn, resulted in a capacity/premium crunch and new entrants seeking to capitalise on the robust rating environment. There is clear value in Cyber products for direct insureds which has driven the rapid growth in the market over the last 10 years, and shows no signs of slowing down. With the market now applying strong discipline to their underwriting of Cyber risk, we believe the class will return to healthy levels of profitability and be a good diversifier from traditional P&C writings. One such way in which underwriting discipline is being applied in practice is through the use of 'Outside-In' data at various stages of the policy life-cycle and the insurance value chain.

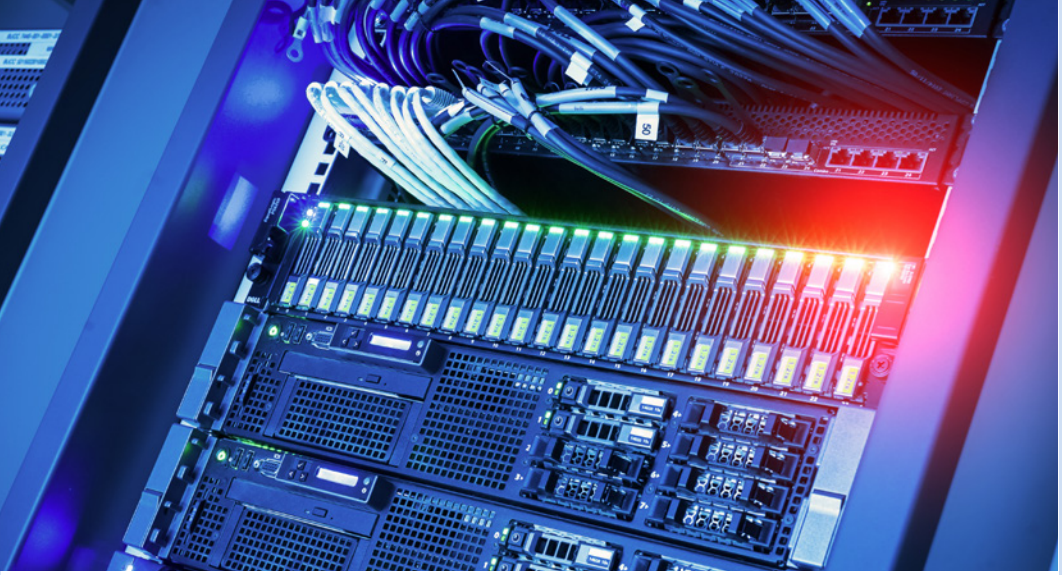

'Outside-In' refers to a collection of data points comprising a company's externally facing IT infrastructure which can be scanned 'from the outside'. This data is often harvested by specialised technology companies in an automated manner and at a large scale, to build a view of the security posture of individual companies. The results of the scans are often then condensed into various buckets of risk factors e.g. Social Engineering, Patching Cadence and Leaked Credentials and an overall score, which is provided as a report to organisations. This helps companies with their own internal risk management and Cyber Security, but is also used to evaluate their trading partners and supply chain risk.

This technology provides clear value to carriers underwriting Cyber policies, providing them with an objective third-party view of an insured's externally facing infrastructure without the need to ask them to fill in a long questionnaire. External scanning data also offers a glimpse at how an insured's security controls are operating in practice, as opposed to how they're designed. Whilst there are certain limitations to the use of Outside-In data, we detail in this report some of the main areas we believe it can positively contribute to the understanding of Cyber risk in the insurance value chain, including:

- Sales & Marketing
- Underwriting & Pricing
- Attack Management— 'Emerging widespread event'
- Pre-Incident Services
- Portfolio Management & Accumulation
- Case Studies in Practice

At Gallagher Re, we have been exploring the potential for this data to help cyber (re) insurers predict claims and comprehend underlying risk quality. We uncover below how this 'Outside-In' technology is currently being deployed by Cyber (re)insurers and analyse case studies of how well it captured past cyber events. Our large scale analysis in Summer 2022 aims to provide (re)insurers with a recipe for which external scanning data ingredients best predict claims, enabling them to offload those that are less valuable and potentially clutter decision-making (false positives).

Whilst external scanning technology has clear limitations, the findings of our initial study have been promising. Providing an attacker's view of the risk, Outside-In data has a key role to play in the future of Cyber (Re)Insurance and shows potential to be predictive of past claims on a portfolio level, but great care needs to be taken to use the data in an appropriate, targeted and appropriate manner.

At Gallagher Re, we have been exploring the potential for this data to help cyber (re)insurers predict claims and comprehend underlying risk quality.

Providing an attacker's view of the risk, Outside in data has a key role to play in the future of Cyber (Re)Insurance and shows potential to be predictive of past claims on a portfolio level, but great care needs to be taken to use the data in an appropriate, targeted and appropriate manner.

Agenda

What is 'Outside-In' Data?

'Outside-In' is a collection of data points that aim to indicate a subject's cyber security maturity without requiring access to the internal network of the subject (e.g. insured).

Outside-In technology vendors aim to help users view a risk from the attacker's perspective. Data is often conveyed as a single score or a selection of factors focused on different areas of risk.

In practice, the data comprises a diverse range of technographic data and firmographic data spanning the technological footprint and attack surface of a company.

Outside-In Technology' — A brief history

Providers of 'Outside-In' data have been around for over a decade, operating across a broad spectrum of use cases. One of the most prominent and successful is offering individual companies a 3rd party view on their cyber defences. This view provided from 'the outside' is critically the same information an external threat actor would see when they first begin reconnaissance of their target. Associated to this is looking at the cyber defences of your core trading partners or supply chain, as attackers often seek to leverage connected entities (predominantly with weaker security countermeasures) as a means of entry to their main target.

The external scan performed by an Outside-In provider may detect potential openings for attackers such as open RDP (Remote Desktop Protocol) ports, unpatched vulnerabilities and poorly configured web services (see later Case Studies for detailed examples). While there is clear value in the information provided by this technology, it has also faced a number of challenges around the credibility and availability of the data, along with the speed in which it can be gathered, cleansed and updated.

Most of the data is collected in an automated manner, scraped by algorithms that scan the entire externally visible IP space and IT infrastructure. Tying back identified assets/ infrastructure to a specific organisation is non-trivial, raising concerns over 'false positives' particularly when scanning larger and multinational companies. The sheer volume of data collected and the associated technology needed to continually scrape it further complicates the issue.

Although many of the previously mentioned challenges still persist, the past five years have seen rapid development in the technology's capability to deliver to the requirements of multiple insurance use-cases:

- Time taken to scan has broadly shrunk from days to minutes, making it viable for underwriting smaller risks or even automating this process;
- Greater historical data has made parameterising models and back testing more credible;
- and, vendors have uncovered ways to discard more false positives, sanitise data and translate technical findings into what they mean for insurance risk.

This evolution of Outside-In technology is akin to the advancements of natural catastrophe models since the early 90s to where they are today, albeit in a more condensed timeframe.

Data Type	What is it?	View on risk	Dynamics to Consider	Summary
Outside-In	Externally available technical and firmographic data aiming to indicate a company's security posture	**How security operates in practice for external facing assets**	**Difficult to Master** (requiring expertise to translate data into insights) **Utility across Value Chain** (from UW to portfolio optimisation and event response)	When used appropriately and proportionately, elements of outside-in data can be used to help predict incidents and subsequent claims
Inside-Out	Data requiring access to an organisation's internal network.	**How security operates in practice within an organisation**	**Uptake requires incentivisation** **Data integration can be automated**	The 'crown jewel' of data for UW and portfolio analysis, but requirement of an inside the network view hampers uptake
Traditional Underwriting	Traditional instruments to manage exposure and reduce risk in UW	**How security controls are designed**	**Can't be entirely replaced by technology** (Provides a view on people and process aspects of security) **Enables proactive response to threat landscape change**	Remains an irreplaceable aspect of understanding Cyber risk

Figure 1: Putting 'Outside-In' data in the context of other understanding of risk enhancing approaches.

The above figure looks at the different views on risk that types of data offer. It is worth briefly mentioning 'Inside-Out' data, which requires access to an organisation's internal network and can provide a view of how security operates in practice inside that network. Similar to telematic devices in cars, a tool (virtual or physical) is planted within internal IT systems that can observe and report on a wide range of security measures, along with the actual behaviour of a company in real-time. Unlike with 'Outside-In' data, the reach of the data observed depends on the scope available to the internal software or device, which can range from restricted browser plugins to full scale, under the hood devices that see almost everything. Technology providers, such as Microsoft with their 'Secure Score', are increasingly making it easier for Insurers to utilise Inside-Out data, by enabling companies who use their software to share security reports in an automated and easily digestible manner. Insurers may even be able to utilise this data throughout a policy lifecycle, providing the opportunity to proactively notify Insureds where their risk profile changes.

Whereas external scanning can visualise external facing vulnerabilities, Inside-Out data provides a complementing view to both Outside-In and questionnaires by aiming to indicate how well an organisation can repel an attack once initial compromise has been made. Inside-Out data can also verify standard security questions asked by UWs such as the use of MultiFactor Authentication (MFA) which cannot easily be determined by Outside-In technology or questionnaires. As a result, Inside-Out data is inherently more insightful when considering readiness of any given firm to respond to an attack.

However, Insurers looking to leverage Inside-Out technology have been met with some resistance as organisations are hesitant to give such wide reaching access and visibility into their internal networks. Gallagher Re anticipates significant application of Inside-Out data for Cyber (re)insurance over the coming years, but challenges around incentivising uptake mean higher barriers to entry for effectively leveraging the technology. The current hard market conditions may have a positive impact on uptake rates, as the potential premium discounts on offer for using Inside-Out technology become relatively more appealing for insureds.

Early adoption by the Cyber Insurance Market

It should come as no surprise that Outside-In providers quickly realised there was a strong use case for their technology within the insurance market as individual organisations sought to transfer Cyber risk. As such, external scanning providers began to actively engage with leading Cyber carriers as early as 2013. The promise of being able to provide underwriters with a view of risk—similar to that of what an external threat actor would see—was enticing and many partnerships were formed between carrier and technology provider.

Figure 2 below shows a summary of external scanning data's history for cyber insurance, along with some of the main players in the space.

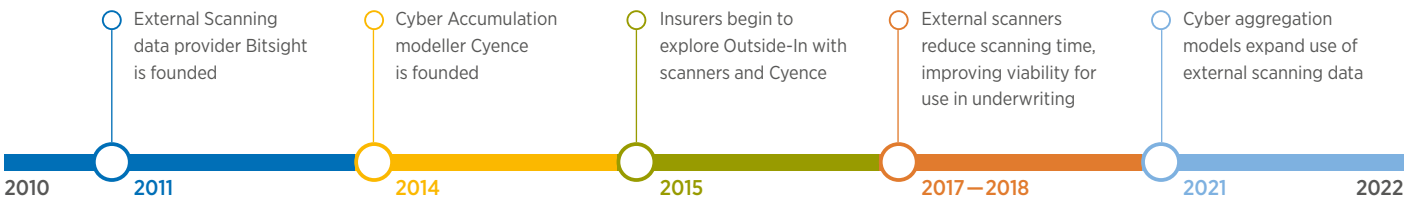

Figure 2: Timeline of external scanning data developments for Cyber Insurance

Whilst the data can't provide a holistic view of Cyber Security posture, it can shed light on elements which traditional questionnaires find hard to capture, focusing on how security controls are 'operating in practice' as opposed to how they're 'designed'. It also provides a regular, year-round view rather than at a single point in time. External scanning data also largely captures different elements of a risk to traditional questionnaires, meaning an outright replacement of questionnaires is unlikely, but Outside-In technology could offer a valuable, complementary view.

External scanning data provides (re)insurers with the attacker's view on risk, largely representing the same pool of data that attackers use in conducting reconnaissance and choosing targets. Many attackers automate their reconnaissance, scanning repeatedly for specific vulnerabilities, by targeting the same vulnerabilities in risk selection and ongoing portfolio validation, (re)insurers can aim to decrease the likelihood of policyholders falling victim to these attacks. Even where this data proves inaccurate—a company appears vulnerable, but has compensating controls in place—an attacker may still choose to target this company based on the perceived vulnerability.

Early test cases of Outside-In data were hampered by some of the previously discussed challenges, including long run times, data accuracy issues for enterprise and SME companies and unsatisfactory numbers of false positives clouding results. Perhaps another issue was the smaller size of portfolios historically, where it was hard to justify the cost associated with a technology solution instead of hiring additional underwriting staff. Furthermore, the overall Cyber insurance market prospered in 2013 through 2018, producing profitable results and little demand for such solutions.

This has taken a turn in recent years, where the dramatic rise in attritional ransomware claims has led to the hardening of the class, significant rate rises and reduced capacity. Underwriters have begun to ask for additional information from insureds and become more selective around which exposures they take on. This desire for additional information (without over-burdensome questionnaires for insureds) is exactly where Outside-In providers can assist.

At the same time, the technology itself has advanced materially since 2013, with issues around runtimes, overall coverage and more detailed, verified information all being addressed. It is for this exact reason, twinned with the desire for more information to underwrite risks that we are witnessing and predicting a continuing uptick in the adoption of this technology for the insurance sector.

We already see this happening in one area of the market which have historically sought to innovate and differentiate themselves: the 'tech-led'. Many of these early adopters were MGAs who saw the opportunity to either leverage or even develop their own in house 'Outside-In' technology, providing an additional view for risk selection and actively monitoring of their overall portfolio of insureds.

In the following section, we'll explore in detail the various current and future applications of 'Outside-In' technology for Cyber carriers, but first want to draw specific attention to the topic of 'Single Points of Failure (SPoFs)'. As mentioned in the executive summary, a key concern for any Insurer's board of directors is the potential scope for systemic events that simultaneously impact many policies. In the natural catastrophe market, these events are well understood, where the aggregating factor is a specific geographical region (e.g. Florida) and portfolios are diversified by writing policies spread around the world. In Cyber, these aggregating factors are 'SPoFs', such as commonly used technologies or services like Amazon Web Services (AWS) or Microsoft Windows Operating Systems (Windows OS). Outside-In expertise has the potential to identify these SPoFs (although not without flaws, discussed later), which has many applications such as understanding accumulations and optimising portfolios to be 'diversified'.

However, external scanning is not yet a trusted and proven technology with a long history of success, and traditional market players remain rightly cautious around relying on the ability of external scanning data to predict claims.

A key concern for any Insurer's board of directors is the potential scope for systemic events that simultaneously impact many policies.

Use of Outside-In Technology through the Insurance Value Chain

External scanning is versatile and the industry has uncovered many potential use cases throughout the insurance lifecycle. As Figure 3 shows, these practical applications aren't solely restricted to underwriting, with the potential to leverage the technology post-bind. We'll explore how insurers are deploying each use case in greater detail below.

Figure 3: Uses for external scanning technology against the Insurance Lifecycle.

There is a significant disparity between how different types of (re)insurers utilise the technology (figure 4 explores this below). Many technology MGAs invested heavily early on and own their 'data value chain' (e.g. Coalition's purchase of Binary Edge and CFC's purchase of Threat Informer, respectively). This means many are capable of doing their own scanning, consolidation of data and integration of that data into processes with limited third-party involvement. Here, they often retain greater 'malleability' of the data and are unencumbered by legacy systems and processes that traditional Insurers grapple with. However, the challenge remains for the technology MGAs as for anyone: even if the data has predictive value for claims, can it be effectively translated into better risk selection and subsequent loss ratio improvements.

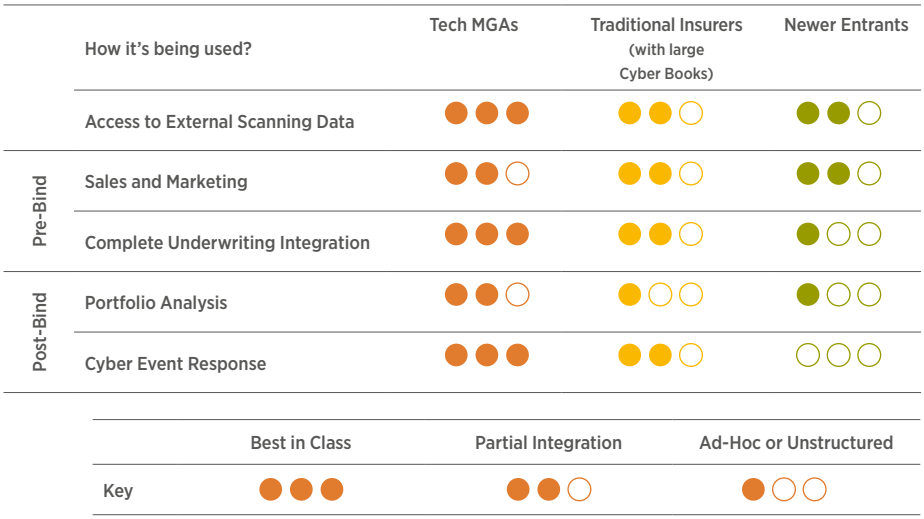

Figure 4: The relative uptake of external scanning data for different use cases by different categories of cyber (re) insurers. This isn't reflective of all (re)insurers in different categories as outliers exist.

Of the vendor landscape, market leaders (such as Bitsight, Security Scorecard and Upguard) are also investing heavily in the application of their data for Insurance and focus on offering traditional insurers more flexibility around how they utilise data. One challenge for insurers is how to translate detailed technical insights into actionable information for decision-making. This requires the expertise to interpret the data and the effective integration of technology into processes to avoid manual dependencies. A new wave of external scanning technology vendors (e.g. KYND, Paladin, Axio, BlackKite) are working to make this easier, aiming to 'lower the barrier' for leveraging Outside-In data by focusing on how that data can be 'translated' into actionable insights for Insurance.

The recent Lloyds Thematic Review encourages the uptake of external scanning technology for multiple use cases across the Cyber Insurance life cycle. Overall, Outside-In data is scalable and offers a different perspective on risk to other technologies available. As such, we anticipate its uptake to become more widespread across the use cases discussed in this paper over the coming years. We explore how the technology can be deployed below:

External Scanning Technology and Company Size

Outside-In technology has different strengths and drawbacks when used to scan companies of different sizes. Whilst certain data points are arguably useful for large entities (e.g. identifying actionable leaked credentials in a timely manner), **the complexity of larger companies renders it difficult to establish a clear picture of company security posture without focusing on targeted external scanning data points.**

Smaller companies with websites usually have simpler footprints meaning scoring is more likely to paint an accurate reflection of security posture. However, hosting providers will often utilise shared and dynamic IP addresses for SMEs, making it difficult for data providers to clearly delineate where one company ends and another begins, risking a reduction in the accuracy of data collected.

Underwriting and Pricing

Perhaps due to the complexity of the data, adoption of various scanning tools and differences in underwriting techniques; no two insurers are alike in how Outside-In data is used for underwriting, but deployments can be summarised by the following categories:

Check and Collect: Utilising external scanning data reports as a check and balance, to compare against own underwriting decision-making without any direct impact to pricing. This enables underwriters to consider whether scanning results generally match up with questionnaire results and provides a window whether those results tell a wildly different story to how security controls appear to be applied in practice. Underwriters may use this comparison to request further information from the insured or challenge questionnaire responses, whilst sidestepping concerns around over-relying on technology.

Key Considerations:

- This 'Check and Collect' approach is arguably the simplest way (re)insurers can introduce external scanning data to underwriting.
- Successful implementations here will ensure results can be gathered as part of the underwriter's workflow without a requirement to log into additional systems.
- The potential external scanning data to add value is limited by ad-hoc integration with underwriting systems and often limited data analysis.

Augmented Underwriting: The semi-automated integration of Outside-In data into the underwriting process. More commonly seen when evaluating SME risks, insurers will often collect data through an API and process it along with questionnaire results to compare a potential insured with risk appetite.

Key Considerations:

- Here, data is often used to identify 'red flags' or require further evaluation under certain parameters. The data may then either directly or indirectly influence pricing decisions.
- Deeper integration of the technology into underwriting processes enables a more proactive approach to translating data into actionable insights for the Insurer.

'Targeted' Deployment': Highlighting practical weaknesses known to be indicators of different cyber events or targets for ransomware crews and other threat actors. Insurers will often reject applications based on these findings, or offer reduced coverage/higher premiums. Perhaps the most common 'targeted' use is a scan for Remote Desktop Protocol configuration, the initial compromise vector in approximately nearly 50% of ransomware attacks.[1]

Key considerations:

- This more 'targeted' use of Outside-In data outlined above may enable insurers to manipulate the dataset to sidestep some of the known limitations e.g. accuracy of overall scores.
- Successful utilisation of external scanning data for 'targeted' deployment will require ongoing access to expertise for translating data into actionable insights and benefit from leveraging historical portfolio data to drive accuracy.

[1] https://www.coveware.com/blog/2021/10/20 ransomware-attacks-continue-as-pressure-mounts

Sales and Marketing

External scanning data can also be used for sales and marketing. Bundling in upfront and ongoing security scans in the premium presents an attraction to policyholders, particularly SMEs, as an extra layer of security. This is the only use case where the quality of the data itself is arguably secondary to its presentation, where translating key findings for a non-technical audience being the key differentiator for vendors to compete. Here, there are two ways by which data is used:

Cyber Health Check Report—provision of a report to a potential insured to support them in developing their understanding of their own Cyber maturity and how they may benefit from Cyber Insurance. This is usually adopted for SME insureds who may lack access to cyber security technology and expertise internally;

Broker Aid—provision of the report to non-technical direct broker with the aim of helping to enrich conversations around cyber insurance, what coverage a potential insured needs and they may benefit from that coverage.

Attack Management

As explored above, the ability to view risk from the attacker's perspective has utility in proactively supporting insureds in responding to unfolding events. Even where compensating controls are in place, the external appearance of weakness may paint a target on an insured.

The Microsoft Exchange attack in 2021 saw potentially thousands of servers compromised as multiple zero-day vulnerabilities for one of the most commonly used services appeared overnight. Multiple threat actors were able to exploit these vulnerabilities, owing in part to the publically scanable nature of the servers. This 'visibility' benefited attackers, but also aided some (re)insurers with access to relevant external scanning tools, by enabling them to determine which insureds were potentially vulnerable to the attack and communicate with them, highlighting the potential exposures and offering good practice advice. For example, the technology MGA Coalition assert they "were able to notify and remediate the vulnerability for 98% of... impacted policyholders within a week of the disclosure."[2] Whilst they suffered claims from the residual 2% of insureds, it's likely that proactive notifications prevented claims and potentially limited the severity of those suffered.

In many cases, the public visibility of a vulnerability is the feature that means a Cyber event has a potential CAT footprint. External scanning data doesn't account for compensating security controls an insured may have in place and a sophisticated insured may appear vulnerable as a honeypot to ensnare would-be attackers. However, adopting Outside-In data in response to Cyber events arguably offers potential for limiting the severity of similar incidents suffered by Insureds or even help some avoid it altogether. Market-leading vendors have begun to offer this capability to (re)insurers following major relevant incidents. The UK's National Cyber Security Centre (NCSC) has also created a similar service (dubbed an 'early warning system') with the potential for integration with Insurers.

In many cases, the public visibility of a vulnerability is the feature that means a Cyber event has a potential CAT footprint.

[2] Cyber Insurance Claims Report – Coalition – H1 2021

Pre-Incident Services

Pre-incident services is the broadest category for use of external scanning data, the most underdeveloped and possibly holds the highest potential to add value. Utilisation of external scanning data in pre-incident services usually refers to the provision of a vendor licence to insureds as a value add, enabling them to manage their own security maturity (often with the insurer retaining access to the licence too). From a (re)insurer's perspective, this potentially provides an additional line of defence, by empowering the insured to better identify vulnerabilities and changes to their risk throughout a policy lifecycle.

The greatest uptake of this technology has been with SME risks as larger entities often already have developed cyber security functions and access third-party scanning tools. External scanning insights can be bucketed with other value add products in a 'CTO as a service' model. This is particularly useful for a portion of SMEs, for whom a 'protection gap' renders it difficult to identify and prohibitive to purchase the right products to ensure security baselines are in place. Whilst all other uses of external scanning data have grown as the market has hardened, uptake of technology for pre-incident services appears to have slowed. This may be partially due to the development and rise in uptake of 'Inside-Out' technology, which provides the insurer and/or insured with a view of their cyber risk 'beyond the firewall'. Here, the provision of security-enhancing 'Inside-Out' services and tools through Insurance can help rebrand an otherwise uninsurable risk to one within appetite and provide the (re)insurer with insights on how their risk exposure is evolving throughout a policy lifecycle.

Portfolio Management

The hard market has incentivised (re)insurers to explore potential tools, beyond firmographic data (largely comprising company size, industry sub-sector and geography) and refreshed questionnaires at renewal, for effectively managing and optimising their portfolios. As a technology that provides a different perspective on Cyber risk when compared with questionnaires, external scanning presents a compelling tool in a (re) insurer's arsenal to evaluate a portfolio against new risk tolerances, or even to make retention decisions. For example, observing the percentage of your portfolio with open Remote Desktop Protocol and combining this with different data points may help indicate your portfolio's exposure to attritional ransomware losses.

Another feature of external scanning utilised by some (re)insurers to evaluate portfolios is its ability to indicate possible third-party dependencies of an Insured, or SPoFs (Single Points of Failure) at a portfolio level. This can also be used for catastrophe modelling and is a focus area for development in vendor aggregation models e.g. cloud outage and systemic ransomware, with CyberCube, Guidewire and AIR utilising SPoF data and Cybercube integrating the technology into its Portfolio Manager in 2021. Here, the technology captures observed connections a company has with others that are visible externally; from ISPs (Internet Service Providers), to cloud services such as AWS (Amazon Web Services).

The allure for (re)insurers here is to move beyond using market share data to determine accumulation risk in portfolios. Whilst the technology's use shows great promise in identifying SPoFs, it's still at an early stage of development, suffering from false positives and limited in ability to determine exactly how a third-party service is being used by an insured. For example, a connection to Amazon Web Services might be recognised, but this doesn't capture the scope or scale of use. This usage might just be a free trial, or another service using AWS. As a result, many vendors still make arbitrary assumptions around the threshold by which a connection observed becomes a dependency for an insured. As a result, we would encourage its use in portfolio optimisation restricted to where a (re) insurer is able to self-determine an appropriate threshold.

External scanning has the potential to add value right throughout the insurance lifecycle. A general trajectory of overall increasing uptake amongst (re)insurers masks the technology's more patchwork usage and often limited integration with existing processes. Validation on Outside-In's ability to deliver against these use cases is promising despite limitations. Whilst early adoption challenges persist—namely, how much emphasis to place on an emerging and partially understood technology—those who've invested in validation and understanding external scanning's propensity to predicting claims early will be well placed to reap the future rewards of stronger risk selection.

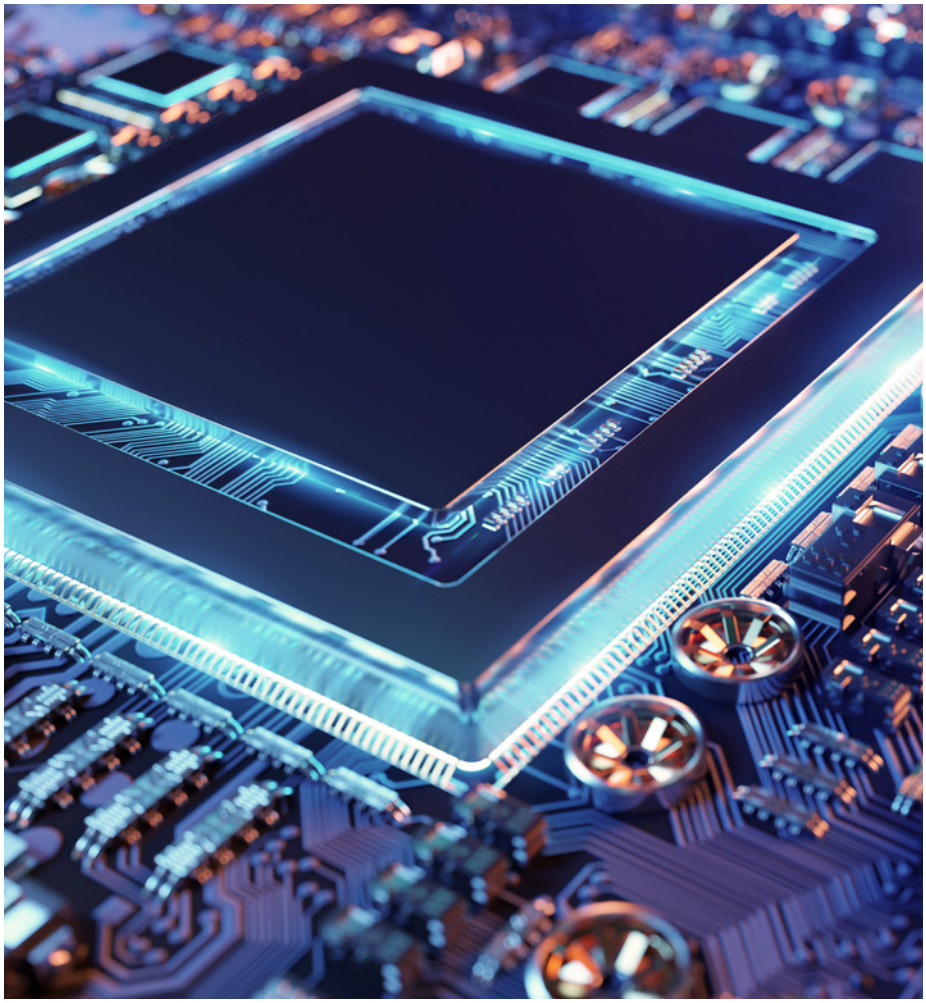

How well does external scanning data predict cyber events and claims?

It's possible for us to discern threat actors are targeting certain features of a company's visible attack surface (e.g. ransomware actors utilising exposed RDP to conduct an initial compromise). However, our lack of comprehension around the dataset's predictive power makes it hard to put external scanning's drawbacks in context and evaluate vendors and their data objectively. This is why understanding the predictive power of Outside-In data to anticipate Cyber Insurance claims is more important than ever to enabling the industry to place reliance on the technology in an appropriate and proportional way.

Many vendors using Outside-In data are working hard to better understand how predictive their models are for different types of Cyber event in any given period. However, most don't have access to the data required to conduct this analysis against claims. Gallagher Re are currently working to better understand how predictive Outside-In data is at identifying cyber incidents and claims. These studies won't offer overnight statistical certainty and the evolving threat landscape will mean the predictive nature of different data points is always shifting, but positive results around the ability of some external scanning data to predict claims has encouraged us to conduct a more detailed study against a larger dataset. We will have the results of this larger study in Summer 2022.

Gallagher Re are currently working to better understand how predictive outside in data is at identifying cyber incidents and claims.

Our lack of comprehension around the dataset's predictive power makes it hard to put external scanning's drawbacks in context and evaluate vendors and their data objectively.

Our Study

Our analysis of the data has been threefold:

1. A comparison of different overall vendor scores;
2. A comparison of different vendor 'risk factors' and their ability to predict claims against a small data set;
3. A comparison of different vendor data to capture individual cyber events.

Five vendors were included in our study, but only four of these returned the data we required to conduct the aforementioned testing. Three of these vendors managed to return data on over 90% of requested companies. Company matching is a significant short-term challenge for the successful utilisation of external scanning data, with most vendors requiring a URL (website) to conduct analysis with accuracy.

The Scores

Many vendors seek to consolidate their data into a single score. When analysing the overall scores of vendors, we found limited consistent view of overall risk. This is likely due to different scoring methodologies between scanners, whose scores focus on solving different problems, like helping a CISO understand their supply-chain risk. For our set of companies, we ranked the overall scores at one point in time, to compare the results between vendors and assess their relative view of risk.

Figure 5 below shows the range of ranked scores for select companies in our portfolio, with a taller vertical bar showing disagreement between vendors. A wide range of scores was observed for most companies, suggesting the scoring methods are significantly different between data providers. Whilst we found limited consistency in the overall relative view of risk, we found that the use of an overall score to predict claims may still add value.

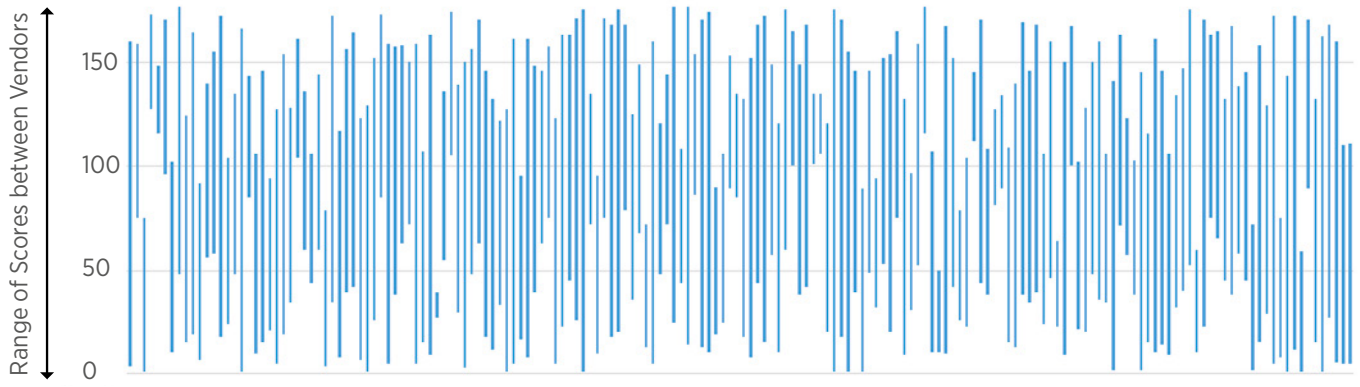

Figure 5: The range of ranked scores for select companies in our portfolio. A taller vertical bar shows disagreement between vendors and each bar shows a specific company. Note the inconsistency.

The Risk Factors

Vendors in our study, and external scanning companies more generally, are broadly utilising the same underlying data points as the ingredients for their scores. However, some scanners have access to additional data points through investment heavy specialisms. For example, we found that data providers which had invested heavily in sinkhole infrastructure had comparatively less false positives when considering which devices are part of a 'Botnet'.[3] Whilst there's limited differentiation in the underlying data points, we found all vendors took different approaches when consolidating these into risk factors or indicators with some suppliers weightings placed on or groupings given to different risk factors. We see these risk factors to be the sweet spot for use for many (re)insurers, providing flexibility and a middle ground between too much data and too restrictive consolidation. We focused on these risk factors for our predictive study.

We used machine learning to ask the question, which of these factors are potentially most predictive of claims? Repeating our method for data breach and ransomware specifically. Overall, we were surprised with how predictive some of the factors we reviewed looked in our results. 'Patching Cadence', regardless of vendor differences in calculation looked to have the potential to predict ransomware claims, whilst others, such as those focused on capturing 'Leaked Credentials' showed more predictive variance between vendors. Whilst some factors showed great potential, others, such as 'Email Security' showed significantly less promise. For anything dark web related, it was difficult to determine if the results we were seeing related to hackers planning an attack, or celebrating one that had already taken place. There were also surprises, where factors and data anticipated to be good indicators of specific types of attacks, showed limited predictivity across overall claims at a portfolio level. This may partially be due to sample size, but our hypothesis here is that whilst some data points are correlated to certain types of events, they may not be with other types of incidents. These are all elements we'll look to investigate with our larger study in 2022.

[3] Sinkhole Infrastructure enables the redirection of traffic on a network through a route of the sinkhole operator's choosing. There are legitimate and malicious applications for sinkholing. In the case of external scanning vendors using sinkhole infrastructure, the aim is to divert requests from a botnet to their command server and identify IP addresses in the botnet. When scanning a company, the vendor can then compare company associated IP addresses with botnet associated IP addresses.

The Case Studies

Beyond taking the portfolio view, we analysed case studies of individual cyber events that had resulted in claims. We compared the hypotheses we had of what should be visible from these attacks externally with the reality of what vendors were able to capture over time. Across our case studies, external scanning vendors were able to capture many of the aspects we'd anticipated of attacks taking place, with a few exceptions. However, our case studies highlighted possibly the greatest limitation of external scanning technology today, that findings at the data point level were often not amplified by vendor algorithms to enable decision-making. With the underlying value was so often obfuscated by noise when the data was consolidated. We explore three of these case studies in further detail below:

CASE STUDY 1: Ransomware attack through RDP compromise

Company Type: Large Corporate
Year: 2020
Event: Targeted Ransomware

	Vendor 1	Vendor 2	Vendor 3	Vendor 4
Event Visible Externally	Yes	Yes (not scanned before October 2020)	Yes (not scanned before March 2020)	Yes
Event visible in Risk Factors	No	No	No	No

A large corporation was hit by a ransomware attack causing encrypted assets and potential data loss within the period of our study. Once an advanced ransomware crew successfully compromises a network, we would expect the large corporation to suffer some impact, with only the most sophisticated security measures able to completely neutralise the attack. The group responsible for this attack is known to use Remote Desktop Protocol (RDP) weaknesses as the entry point for almost all attacks. As these weaknesses are theoretically visible from the outside, we analysed whether vendors effectively captured this exposure.

Hypothesis

Web malware is captured by external scanning in the form of weak security controls on organisations website, which either made the attack possible or are resulting from attacker actions.

Reality

The event was visible from the outside for three of the four vendors analysed. However, without further details, decision makers wouldn't be able to understand the nature of the event and take remedial actions.

All vendors were different in their categorisations of Web Security.

Two vendor's overall scores dipped within weeks after the attack became public. However, this was likely because these vendors artificially reduce overall scores after an event has been observed for a pre-defined period of time and not due to changes in the underlying data.

Vendors aren't always able to capture compensating controls an organisation has in place to harden RDP, but they were consistent in their visibility of RDP exposure across our portfolio.

When observing similar events in our sample, we noted that RDP exposure was binary; regardless of the number of instances of exposed RDP, attackers seemed able to exploit the weakness. The limited ability of vendors to highlight these weaknesses in risk factors and overall scores is largely because RDP is one of hundreds of data points.

CASE STUDY 2: Credit Card Skimming Malware

Company Type: Midsize Retailer
Year: 2020
Event: Web Malware

	Vendor 1	Vendor 2	Vendor 3	Vendor 4
Event Visible Externally	Yes	No	Yes (not scanned before March 2020)	No
Event visible in Risk Factors	Partially	No	Partially	No

A midsized consumer goods and e-commerce company was hit by website credit card skimming malware during the study period. This type of malware usually exploits vulnerabilities in a company's web platform to embed malware. We analysed the data to see which (if any) web platform vulnerabilities were visible from the outside.

Hypothesis

Exposed database and nature of the data exposed is visible externally for the period of misconfiguration.

This data is captured by vendors and highlighted in risk factors enabling decisions to be taken with knowledge of the vulnerability's existence.

Reality

Whilst all vendors had the potential to capture the incident in underlying data, only one was able to amplify this finding and make it visible to ourselves.

This vendor also reported on the specific database vulnerable and the types of data exposed and promptly reported the incident as fixed.

Whilst all vendors are uniform in consolidating patching cadence into a single risk factor, they all take different approaches with regards to Web Security, which made like-for-like comparisons difficult. Although the event was theoretically visible in two of the vendor results, it would have been difficult to translate what the vendor findings meant in the context of the incident without significant further analysis for the most technical reader. Like many of the case studies we reviewed, visibility of this event suffered from the limitations of the vendor algorithms in overcoming the noise of other data points and amplifying the issues associated with the attack.

Case Study 3: Exposed Cloud Service

Company Type: Large Corporate
Year: 2019/2020
Event: Data Leak

	Vendor 1	Vendor 2	Vendor 3	Vendor 4
Event Visible Externally	No	No	Yes	No
Event visible in Risk Factors	No	No	Yes	No

Cloud-based storage often requires users to manually configure security settings so that data is secured. A large corporation was hit by a data leak owing to a poorly configured wordpress database within the period of our study. In this instance, the data was theoretically publically exposed and visible to all malicious and non-malicious scanners. We analysed the data to understand if Outside-In providers picked this up.

Hypothesis

Attacker used exposed RDP (Remote Desktop Protocol) to conduct initial compromise of victim. This exposure is captured by vendors and highlighted in high level risk factors.

Reality

All vendors successfully captured the exposed RDP for the company which suffered the attack.

No vendor amplified this finding at the risk factor level in a material way. The only movement observed at the overall score level was after the attack became public.

The challenge for other vendors to capture this risk was further compounded by the size of the organisation in question. The complexity of large corporates makes it more challenging for external scanners to remove false positives, identify findings and report them with minimal manual oversight. For this reason, questionnaires and/or 'behind the firewall' information may be preferable when evaluating large corporates.

Conclusion

Throughout this report, we've observed clear value that 'Outside-In' technology can bring to the insurance value chain, with multiple use cases where this data is already being used to enhance underwriting and visibility of Cyber risk by carriers. One such success story is in post-event attack management; where carriers can immediately scan their entire in force portfolio through a provider for a specific external facing vulnerability, identifying and contacting any insureds that appear exposed. This proactive risk management should reduce the potential impact from systemic Cyber events where attackers exploit external facing vulnerabilities, adding clear, tangible value to insureds and carriers alike. In the future, external scanning data may also help carriers in estimating their potential losses following a widespread event when combined with aggregation modelling, much like the major natural catastrophe models provide an event ID after a major event to run against a carrier's exposure. We anticipate that external scanning's uptake in Insurance will increase over the coming years, as the technology's potential across more use cases identified in this report are fully realised throughout the insurance value chain.

To assist in the development of this technology and improve the level of confidence in providers of 'Outside-In' data, Gallagher Re have already conducted a preliminary study of 250 companies and several historical events to develop an understanding of which factors are most important in predicting future insurance claims. Whilst we found limited convergence in the overall relative view of risk between vendors (by their aggregated 'Cyber Risk Score'), there appeared to be value in many of the underlying risk factors and data that are combined to generate the overall score.

Most 'Outside-In' providers are utilising largely similar underlying data sets, but the real differentiator is how they consolidate and translate that data into actionable insights on an organisation's Cyber Security posture for decision makers. As such, one of the most critical points that our preliminary study highlighted is the ability for providers to identify and strip out false positives with minimal (ideally no) manual intervention. Another major challenge for providers is locating and amplifying important findings in an automated manner to highlight the high risk vulnerabilities or exposures (such as exposed cloud services) that hold the greatest likelihood of an Insured suffering a claim. The information is often there and available, but model weighting by vendors or sheer volume of data may obscure these findings from insurance practitioners. To help combat these challenges, Gallagher Re are conducting a more in-depth study in H1 2022 across our entire industry exposure and claims database to properly identify predictive factors and share these with the market.

Despite our overall optimism, 'Outside-In' data is not without its challenges and will not be a silver bullet for understanding the Cyber risk of an insured. Rather, we believe it will provide a complementary source of information that will enable underwriters to better understand the risk they are insuring and supplement UW questionnaires. Attackers are always going to use all the information available to them to launch malicious attacks and successful (re)insurers looking to stay one step ahead must leverage that same information. External scanning has many limitations and must be deployed in an appropriate and proportional way for success. However, as a powerful tool in a (re)insurers armoury, enabling the industry to leverage an attacker's view to understand evolving exposure and embolden policyholders to bolster their defences, 'Outside-In' is here to stay.

Key Takeaways

- External scanning data provides a (re)insurer with 'the attacker's view', comprising the same information many threat actors use to select and compromise targets.
- Scanning is usually quick, scalable and flexible with potential for use across the insurance value chain. It also provides a complementing view to questionnaires with a partial view of how an insured's security controls are operating in practice as opposed to how they're designed.
- External scanning has limitations. Data accuracy is constrained by false positives, infrequency of updates and sheer data volume; whilst a reliance on individual scores appears to obfuscate the value in underlying data.
- Whilst external scanning data can usually pinpoint how a victim is compromised, it struggles to amplify critical findings in reporting and translate what these findings mean in a digestible format for (re)insurers.
- Overall, our preliminary analysis has indicated the technology has the potential to add value to our understanding of Insured Cyber risk when used in an appropriate and proportional manner.

Authors

Justyna Pikinska
Head of Specialty Analytics

E: Justyna_Pikinska@gallagherre.com

Ed Pocock
Senior Cyber Security Consultant

E: Ed_Pocock@gallagherre.com

Michael Georgiou
Senior Cyber Actuary

E: Michael_Georgiou@GallagherRe.com

Gallagher Re